

Social physics

Marko Jusup^{a,b}, Petter Holme^a, Kiyoshi Kanazawa^{c,d}, Misako Takayasu^a,
Ivan Romić^{a,e}, Zhen Wang^f, Sunčana Geček^g, Tomislav Lipič^h, Boris
Podobnik^{i,j,k,l}, Lin Wang^m, Wei Luoⁿ, Tin Klanjšček^g, Jingfang Fan^{o,p},
Stefano Boccaletti^{q,r,s}, Matjaž Perc^{t,u,v,*}

^a*Tokyo Tech World Hub Research Initiative (WRHI), Institute of Innovative Research,
Tokyo Institute of Technology, Tokyo 152-8550, Japan*

^b*Department of Physics, University of Rijeka, HR-51000 Rijeka, Croatia*

^c*Faculty of Engineering, Information, and Systems, The University of Tsukuba,
Tennodai, Tsukuba 305-8573, Japan*

^d*Japan Science and Technology Agency, PRESTO, Kawaguchi, Saitama 332-0012, Japan*

^e*Statistics and Mathematics College, Yunnan University of Finance and Economics,
Kunming 650221, China*

^f*School of Artificial Intelligence, Optics, and Electronics (iOPEN),
Northwestern Polytechnical University, Xi'an 710072, China*

^g*Division for Marine and Environmental Research, Ruđer Bošković Institute,
HR-10002 Zagreb, Croatia*

^h*Division of Electronics, Ruđer Bošković Institute, HR-10002 Zagreb, Croatia*

ⁱ*Faculty of Civil Engineering, University of Rijeka, 51000 Rijeka, Croatia*

^j*Faculty of Information Studies in Novo Mesto, 8000 Novo Mesto, Slovenia*

^k*Zagreb School of Economics and Management, 10000 Zagreb, Croatia*

^l*Luxembourg School of Business, 1450 Luxembourg*

^m*Department of Genetics, University of Cambridge,
Cambridge CB2 3EH, United Kingdom*

ⁿ*Department of Geography, Faculty of Arts & Social Sciences National,
University of Singapore, Singapore 117570, Singapore*

^o*School of Systems Science, Beijing Normal University, Beijing 100875, China*

^p*Potsdam Institute for Climate Impact Research (PIK), 14412 Potsdam, Germany*

^q*CNR—Institute for Complex Systems, 50019 Florence, Italy*

^r*Universidad Rey Juan Carlos, 28933 Móstoles, Madrid, Spain*

^s*Moscow Institute of Physics and Technology, National Research University,
141701 Moscow Region, Russia*

^t*Faculty of Natural Sciences and Mathematics, University of Maribor,
2000 Maribor, Slovenia*

^u*Complexity Science Hub Vienna, 1080 Vienna, Austria*

^v*Department of Medical Research, China Medical University Hospital,
China Medical University, Taichung 404332, Taiwan*

*matjaz.perc@gmail.com

Abstract

Recent decades have seen a rise in the use of physics methods to study different societal phenomena. This development has been due to physicists venturing outside of their traditional domains of interest, but also due to scientists from other disciplines taking from physics the methods that have proven so successful throughout the 19th and the 20th century. Here we dub this field ‘social physics’ and pay our respect to intellectual mavericks who nurtured it to maturity. We do so by reviewing the current state of the art. Starting with a set of topics that are at the heart of modern human societies, we review research dedicated to urban development and traffic, the functioning of financial markets, cooperation as the basis for our evolutionary success, the structure of social networks, and the integration of intelligent machines into these networks. We then shift our attention to a set of topics that explore potential threats to society. These include criminal behaviour, large-scale migrations, epidemics, environmental challenges, and climate change. We end the coverage of each topic with promising directions for future research. Based on this, we conclude that the future for social physics is bright. Physicists studying societal phenomena are no longer a curiosity, but rather a force to be reckoned with. Notwithstanding, it remains of the utmost importance that we continue to foster constructive dialogue and mutual respect at the interfaces of different scientific disciplines.

Keywords: multidisciplinary, thermodynamics, statistical physics, human behaviour, sustainability

Contents

1	Prologue: The physical roots of multidisciplinary	5
2	Urban dynamics	8
2.1	The definition of a city	8
2.2	Size of cities	10
2.3	City growth	11
2.4	Networks within and of cities	12
2.5	Segregation	15
2.6	Scaling theory of cities	16

2.7	Human mobility	17
2.8	Trajectory analysis	18
2.9	Navigability	19
2.10	Future outlook	20
3	Traffic flows	20
3.1	Observed phenomena	21
3.2	Traffic-flow models	23
3.3	Pedestrian traffic flows	29
3.4	Future outlook	29
4	Econophysics	31
4.1	The advent of econophysics	31
4.2	Agent-based modelling: The dealer model	33
4.3	Time-series modelling: The PUCK model	36
4.4	Order-book modelling: Financial Brownian motion	39
4.5	A kinetic approach to financial market microstructure	42
4.6	Solving the microscopic model via kinetic theory	49
4.7	Consistency between theory and data	54
4.8	Future outlook: towards market ecology	56
5	Cooperation	57
5.1	Social dilemmas	58
5.2	Cooperation in networks with pairwise interactions	62
5.3	Cooperation in multilayer networks	66
5.4	Cooperation in temporal networks	69
5.5	Cooperation in networks with higher-order interactions	74
5.6	Empirical facts about human cooperation	78
5.7	Future outlook	83
6	Networks and communities	85
6.1	Community detection: contexts and methods	86
6.2	Introducing stochastic block models	88
6.3	Defining stochastic block models	91
6.4	Statistical inference of communities	95
6.5	Future outlook	102

7	Human-machine networks	104
7.1	Literature walkthrough	109
7.2	Fundamentals of artificial intelligence	112
7.3	Deep learning	117
7.4	Learning to learn	120
7.5	AI agents for promoting cooperation	122
7.6	Future outlook	127
8	Criminology	128
8.1	The broken windows theory	128
8.2	Crime hotspots	131
8.3	Crime as a self-exciting point process	135
8.4	Social dilemmas of crime	136
8.5	Criminal networks	142
8.6	Rehabilitation and recidivism	149
8.7	Rosy outlooks for less crime	151
9	Migrations	153
9.1	Tolerance	157
9.2	Integration and culture	161
9.3	Populism and polarisation	165
9.4	Future outlook	170
10	Contagion phenomena	171
10.1	Data-driven metapopulation models	173
10.2	Digital epidemiology	176
10.3	Analytical results from metapopulation models	181
10.4	Bayesian inference	191
10.5	Challenges and future work	196
11	Environment	199
11.1	Pollution	200
11.2	Physical habitat loss	210
11.3	Tipping points	212
11.4	Future outlook	219

12 Global climate change	223
12.1 Modelling the climate system	227
12.2 Climate networks	233
12.3 Impact of Rossby waves on air pollution	241
12.4 Critical phenomena in the climate system	244
12.5 Future outlook	251
13 Epilogue: Keeping the dialogue open	251

1. Prologue: The physical roots of multidisciplinary

The present text is perhaps best described as a journey through what has become an extremely active and diverse research field known under the umbrella term *social physics*. Before an interested reader embarks on this journey with us, it is only fair to inform them of our motivation and rationale. Doing so, at the very least, requires (i) defining what social physics means (to us), (ii) outlining the case for its importance, and (iii) elaborating the underlying line of thought that connects the topics covered henceforward.

The methods of probability and statistics first flourished among social scientists who sought quantitative regularities revealing the inner workings of society [1]. This inspired the founders of statistical physics in the 19th century to move away from Newtonian determinism and embrace a probabilistic description of ideal gases. Today, however, physicists are completing the circle by applying physical methods (oftentimes those of statistical physics) to quantify social phenomena [1]. For our purposes here, we decide to adopt a broad, operational definition of social physics. Specifically, social physics is a collection of active research topics aiming to resolve societal problems to which scientists with formal training in physics have contributed and continue to contribute substantially. Although the precision and rigour of such a definition may be questioned, we are in a good company when relying on what physicists actually do to define (an aspect of) physics [2]. We also believe that being inclusive and practical about what constitutes social physics makes us appreciate more the broad position of physics in modern science.

The twentieth century has often been called “a century of physics” [3], and for good reasons too. The scientific method as practised by physicists achieved enormous success on all scales of reality. On the small end of things, quantum electrodynamics as the theory of the interaction between light and matter has been tested to within ten parts per billion (i.e., 10^{-8}) [4] by

examining if the dimensionless magnetic moment of the electron relates to the fine structure constant as predicted. On the large end of things, general relativity as the prevailing theory of gravitation has been tested, among others, by putting satellites in space, measuring the geodetic effect with an error of 0.2% and the frame dragging effect caused by Earth's rotation with an error of about 19% [5] relative to predictions. The latter error, incidentally, amounts to 37mas which is best put into perspective by the words of investigator Francis Everitt that 1mas "is the width of a human hair seen at the distance of 10 miles." What is important for us is that these enormous successes of physics have caught the attention of scientists from other disciplines, and have led to attempts to generate similar successes using physics-like quantitative methods. This is explicitly admitted by some disciplines; in ecology, for example, metabolic theory [6, 7, 8, 9] and mechanistic niche modelling [10, 11] draw heavily from thermodynamics. More generally, though, the adoption of physics-like quantitative methods is best reflected in the proliferation of physical and mathematical modelling in disciplines as diverse as epidemiology [12, 13], virology [14, 15], neuroscience [16, 17], medicine [18, 19], psychology [20, 21], sociology [22, 23], and countless others. But if the ultimate goal is to replicate the success of physics, whom better to call for help than physicists themselves? All this explains to a decent degree why physics is at the roots of the modern shift to multidisciplinary, and why physicists publish more multidisciplinary physics articles than articles in physics journals [2].

Although the success of physics over the past hundred year or so is impossible to dispute, signs of a progress slowdown [24, 25] and dissatisfaction with fundamentals [26, 27] have been brewing, consequently setting in motion multiple searches for 'new physics' [28, 29]. Until such physics is found, however, many a young physicist may seek to employ their strong quantitative skills elsewhere. A well-known example is attempts by physicists to, both through research in academia [30] and practice on Wall Street, enrich the world of finance.

The described state of affairs places physics squarely at the roots of multidisciplinary. It is then hardly surprising that much of the multidisciplinary work conducted by physicists aims at resolving societal problems. We call that work social physics, believing that otherwise we would be diminishing the rightful role of physics in today's multidisciplinary movement. Provided the reader is willing to agree with us or, at least, give us the benefit of the doubt, it becomes glaringly obvious that the scope of social physics is enor-

mous. How did we go about narrowing down a relevant set of topics for the present text?

Our focus was twofold, asking what enables or constitutes the modern way of living and what perturbs or threatens it. The majority of human population now lives in cities [31] because of more healthcare, education, and employment opportunities. Despite their advantages, cities suffer from many problems, traffic being among the more acute ones [32]. We therefore started by overviewing the contributions of physicists to research in *urban dynamics* and *traffic flows*. The prosperity of cities is in many ways tied to markets, and over the past two decades financial markets in particular had an enormous impact on urban life, innovation, and planning [33]. This warrants a better understanding of financial markets, which is the aim of the chapter on *econophysics*. Life in cities, and civilised life in general, is based on widespread cooperativeness. The evolution of *cooperation* accordingly deserves a chapter of its own, even more so given that this is a research domain in which physicists have been especially active [34]. Human population is furthermore organised in social networks, whose structure is entwined not only with the evolutionary dynamics of cooperation, but also many other dynamical processes of societal relevance [35]. Probing network structure and their separation into *communities* could therefore not be overlooked. An important realisation in this context is that computers increasingly take part in shaping social networks, especially so with the advent of human-like artificial intelligence. The present state of affairs and current technological trends, in fact, necessitate a candid discussion about *human-machine networks*.

Among phenomena that perturb or threaten the modern way of living, crime is a conspicuous one, with impacts at large, societal [36] and small, community [37] scales. Interestingly, as the chapter on *criminology* will demonstrate, both the evolutionary dynamics of cooperation and network-structure analyses prove useful in gaining insights into crime fighting and criminal organisations. In contrast to crime, *migrations* per se come with positive effects, such as helping to alleviate labour-force deficits and age-structure imbalances in ageing populations [38], but there are many caveats. Developing countries, for example, were supposed to receive demographic dividends from their favourable workforce-to-dependants ratios, but substantial value has been lost to brain drain, that is, emigration of highly educated young adults [39]. Much more consequential is when population displacements are triggered by environmental shifts or geopolitical instabilities; on the one hand, people losing homes is a humanitarian crisis, while on the other

hand, countries quickly absorbing sizeable immigration fuels nationalism and xenophobia [40, 41]. Movements of people across large distances, especially at a fast pace of today, make humankind vulnerable to contagions [42]. The chapter on *contagion phenomena* covers disease transmissions on global and local scales, and overviews the budding subfield of digital epidemiology. Before closing off this review, we shift the focus to natural surroundings that support human life in the first place. The chapter on *environment* puts emphasis on the proliferation of chemicals whose effects, particularly synergistic ones, remain partly understood at best [43]. Even more critical is *global climate change* [44, 45, 46]. Somewhat surprising in this context is the use of network science to unravel the intricacies of Earth’s climate system. This only goes to show how versatile physics and its methods are, which we hope will inspire physicists to further nurture multidisciplinary, as well as scientists from other disciplines to maintain a dialogue with physicists when resorting to quantitative approaches and tools.

2. Urban dynamics

Cities are archetypal examples of complex systems [47]. They are, to some extent, self-organised, in other aspects planned. They need hierarchical and interdependent distribution systems. They exist across an extraordinary range of scales. They interact in complex networks, they consist of complex networks, and they are built by people interacting through complex networks. Many aspects of urban science do not fit the methodologies of physicists. This section will discuss the current state of urban science [48], in particular the topics that interest physicists [49].

2.1. The definition of a city

Assume that we know the locations of all humans, at all times, and all buildings and infrastructures, then how can we decide what a city is? This is far from an easy question, and one soon realises that any simple solution will not overlap perfectly with the existing conventions. Most of the research that we present here use an administrative definition of a city. We will, however, look at attempts to define cities from the population distribution.

Even if one starts from population density data, one can usually be completely independent of pre-defined administrative borders. For example, Ref. [50] uses the population of the most fine-grained subdivision of

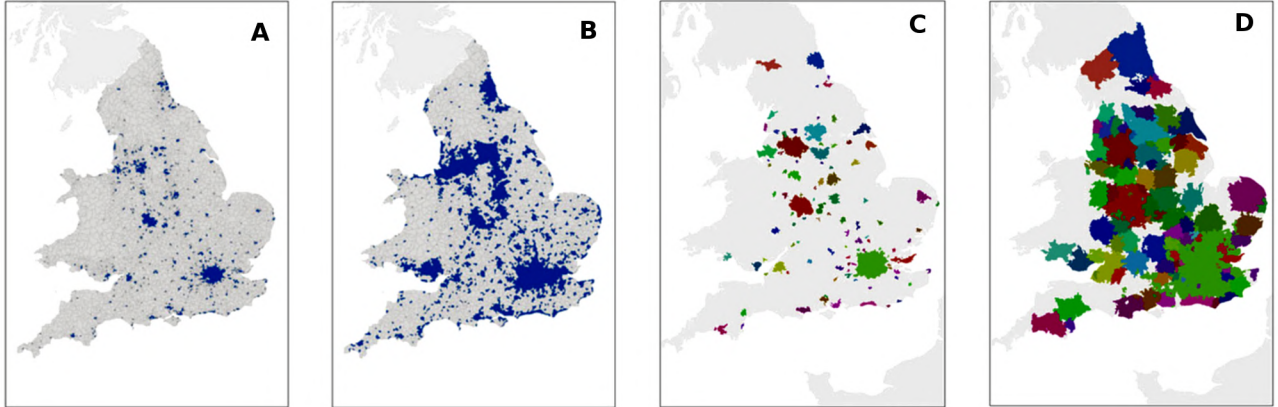


Figure 1: Two methods of defining cities in England. **A, B**, Results of the density thresholding method with thresholds $\rho_0 = 24 \text{ ha}^{-1}$ and $\rho_0 = 2 \text{ ha}^{-1}$, respectively. **C, D**, Results of the commuting-based method with minimum populations of 50,000 and commuting flow thresholds 40 % and 5 %, respectively.

Source: Reprinted figure from Ref. [50] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

the United Kingdom (wards) for this purpose. The actual population count comes from the home address reported in a census.

One idea for the definition of a city, akin to percolation theory [51], is to identify regions with a population density ρ above a particular threshold ρ_0 . Ref. [50] defines a city by requiring that $\rho_k > \rho_0$, for all wards k at the boundary. If wards belonging to the city surround a ward k , we consider it a part of the city even if $\rho_k < \rho_0$. See Fig. 1 for the results of this algorithm when applied to the population density of England.

Another approach to defining cities is to group smaller divisions with larger ones if at least a fraction of the population of the smaller commutes to the larger, and there is not a larger division that attracts even more commuters. In order for this approach to give sensible results, one needs to start from some seed regions more populous than a given threshold. From such a seed region, the algorithm of Ref. [50] recursively adds wards to a region that at least a certain fraction of people commute to. If there is more than one region that draws a fraction of commuters above the threshold, then the ward is added to the region attracting most commuters. When such a recursive procedure has converged, one is left with a city pattern such as in Fig. 1C, D.

2.2. Size of cities

George Kingsley Zipf noted in his 1949 *The Principle of Least Effort* [52] that city sizes m follow a power-law distribution

$$P(m) = \text{const.} \times m^{-\nu}. \quad (1)$$

Zipf found the exponent ν to be one and assumed that it was a universal value, but more recent studies argue that different regions of the world have different ν [53, 54, 55].

There is a vast number of mechanisms proposed for Zipf's law of city sizes, starting from Zipf himself [52]. We will mention a few from the physics literature. First, Zanette and Manrubia suggested a multiplication-diffusion mechanism model operating on a square grid [53]. From an initially even distribution of population density, the population m at a random site i is updated as

$$m_i(t) = \begin{cases} (1-q)p^{-1} & \text{with probability } p \\ q^{-1}(p-1) & \text{otherwise} \end{cases} \quad (2)$$

Then a fraction α of the population is redistributed to the surrounding cells. This model produces emergent power-laws in agreement with Zipf, for a broad range of parameter values.

Ref. [56] proposes a model in which the arrival w_a and departure rates w_d from a city of size m depend on m according to

$$w_a = \frac{m^2}{m_0} + m \quad (3a)$$

$$w_d = e^{1/m^*} \left[\frac{m^2}{m_0} + a \right] \quad (3b)$$

where m^* , m_0 , and a are parameters. These rules are repeatedly applied in combination with a growth of the number of cities (by occasionally adding cities of population one). This gives, for some parameter range, an emergent city-size distribution of

$$\nu = \frac{1}{1 + (a-1)m_0}. \quad (4)$$

In both of the above models, there is an element of 'rich-gets-richer' (often called 'Gibrat principle' [57], sometimes the 'Matthew effect' [58], or 'cumulative advantage' [59]) that larger cities manage to attract more people




Figure 2: Growth of a modelled city. Shown is an artificial city plan generated by the Manrubia-Zanette-Solé model [63]. The colour bar encodes a timeline for the development of the city. The model reproduces many features of real cities, but there are also striking differences compared to reality; in the model, for example, the oldest regions are completely embedded in newer ones.

Source: Reprinted figure from Ref. [63].

and thus grow faster than smaller cities. Thus, many authors cite Herbert Simon’s model for emergent power-law distribution [57] as an explanation for city size distributions. This mechanism has been revived and adapted for city growth in the relatively recent economics literature [60]. Indeed, out of all mechanisms generating power-laws [61], the models specifically trying to explain city growth, that we are aware of, all seem to incorporate a rich-gets-richer mechanism. There is also some direct empirical evidence for a rich-gets-richer growth of cities [62].

2.3. City growth

Another issue about cities that has interested physicists is the spatial growth of cities. Indeed, the Zanette-Manrubia model has also been proposed as a model for the spatial growth of cities [63]. This is maybe not so surprising because, in the spirit of self-similarity, the population distribution within a city could be similar to that of a region containing many cities. Fig. 2 shows one example of a result of this model. Although the model manages to reproduce many features of real city growth, one immediately

spot discrepancies when comparing the model output to real data. The most striking difference is that the oldest regions of the Zanette-Manrubia are completely embedded in newer built environments. However, in real data, they could border non-built land-use types.

In addition to reaction-diffusion type models of city growth, some models are somewhat similar to diffusion-limited aggregation (DLA) [64]. For example, Ref. [65] follows a Markov random field framework, but adds many rules from the urban planning literature or the authors' observations. This model could be coupled with geographic data or similar to improve its predictive power. Another influential paper motivated by DLA, or rather its weaknesses, to model city growth, is Ref. [66]. In this model, the nodes are successively added to the cluster (representing a city) with a logarithmically decaying probability of the distance to occupied areas.

Finally, we note that predicting city growth patterns does not only interest physicists. For example, see Ref. [67] for a recent model by geographers of the co-evolution of land-use and population density.

2.4. Networks within and of cities

As already alluded to, it is not straightforward to demarcate cities from their surrounding. Therefore many of the principles that relate different cities also apply to the organisation of cities themselves. Geographers had invented simple models to explain existing patterns and determine the optimal spatial networks long before physicists turned to this problem (for example, see Fig. 3). We recommend reading Ref. [69], which is an almost 50 years old textbook but will feel very modern for anyone working on spatial, temporal, or higher-order network structures, or the modelling of complex socioeconomic systems.

One influential early model for the network of cities was Walter Christaller's 1933 'central place theory' [70]. It assumes an underlying featureless landscape of uniformly distributed resources. In such a scenario, larger cities would primarily organise in a hexagonal lattice. Secondary, smaller cities would fill the gaps around the central places, and so on. The economist August Lösch derived a more flexible and more economics-favoured location theory than Christaller's in his 1940 *The Economics of Location* [71]. Lösch also concludes that in a structureless world, human settlements would be organised into a hexagonal pattern. Furthermore, cities would have a fat-tailed size distribution [72], although deriving Zipf's law was not an explicit goal of Lösch.




Figure 3: A figure from a 1963 essay by the Greek-American architect C. A. Doxiadis to argue that for a city to grow without losing its vital functions, its centre needs to be supported by a network of subcenters—“to create a new network of lines of transportation and communication which do not lead towards the centre of existing cities but towards completely new nodal points” [68].

Source: Courtesy of Constantinos A. Doxiadis Archives.

© Constantinos and Emma Doxiadis Foundation.

Physicists have spent more effort trying to understanding the evolution of the networks within cities [74, 75] than networks of cities. One example is Barthélemy and Flammini’s model of the growth of street patterns [73]. This model works by adding ‘centres’ that are then connected according to the following rule. Say that A and B are neighbouring centres, and M is a tip of a nearby road under construction. The road will grow from the tip M in the direction of the vector

$$\frac{\overrightarrow{MA}}{|MA|} + \frac{\overrightarrow{MB}}{|MB|}, \quad (5)$$




Figure 4: Street patterns generated by an algorithm. Panels correspond to time progression in the model: (a) $t = 100$; (b) $t = 500$; (c) $t = 2000$; and (d) $t = 4000$. Initially, constructed road form a tree-like structure, but loops start appearing as the density of roads increases.

Source: Reprinted figure from Ref. [73].

such that the cumulative distance from the centres A and B to the road network is minimised (Fig. 4). When the road M reaches a point on the line between A and B , a straight road segment is added from A to B . For further details about this model, see Ref. [73].

Other models of spatial networks typically also operate by successively adding points and connecting these to the existing network [76]. A more general model of spatial growth that could work as a model for road networks is Gastner and Newman’s model in Ref. [77]. This algorithm associates a cost to all the links that is proportional to

$$c_{ij} = \lambda\sqrt{N}d_{ij} + (1 - \lambda), \quad (6)$$

where d_{ij} is the distance between points i and j , and λ is a parameter governing the relative cost of the distance of the link to its existence. Then the algorithm seeks a set of links E that minimises

$$\sum_{(i,j) \in E} d_{ij} \text{ given } \sum_{(i,j) \in E} c_{ij} < C \quad (7)$$




Figure 5: Snapshot of the converged state of Schelling’s segregation model. Symbols ‘O’ or ‘#’ correspond to individuals of two different ethnicities occupying cells in an $L \times L$ square grid. Convergence is achieved by repeatedly following a simple rule such that individuals who have more than a fraction b of neighbours belonging to the other ethnicity move to a another, randomly selected empty cell.

Source: Reprinted figure from Ref. [79].

for a parameter C representing the total budget of the project. If λ is large, the networks become more like urban infrastructures, otherwise the networks are rather like airline maps. The Gastner-Newman model is similar to Fabrikant-Koutsoupias-Papadimitriou model [78] of Internet evolution in the sense that it balances the cost of physical links and the presence of the link in the network.

2.5. Segregation

Studies of segregation by physicists is essentially equal to studies of the model by Schelling [79]. Just one glance at Schelling’s paper should be enough for the reader to understand why this particular model is popular among physicists (Fig. 5). The background of the model was the racial tensions in the 1960s USA in general and residential segregation in particular. Schelling used the model to argue that even if people are mostly tolerant of living close to others of another ethnicity, spatial constraints accentuate segregation. The model works as follows:

1. Consider an $L \times L$ square grid in which every cell can be empty or

occupied by a resident of one of two ethnicities (‘ \circ ’ or ‘ $\#$ ’ in Fig. 5).

2. Initially, distribute $\lfloor fL^2/2 \rfloor$ \circ s and equally many $\#$ s randomly on the grid, where $0 < f < 1$ quantifies occupancy.
3. Update the configuration by picking an occupied square and, if this is surrounded by more than a fraction b of the opposite ethnicity, move it to a random unoccupied square. Schelling considered the eight nearest neighbours.
4. Repeat the previous step until all occupied squares are below the threshold. If f is sufficiently small, the procedure will converge. Otherwise, the problem is ill-defined.

Essentially, the final level of segregation, measured by the average fraction of neighbours of the same ethnicity, will be far from the threshold, and this increases non-linearly with both the threshold b and occupancy f .

There are many papers in the physics literature dealing with this model, typically without interpreting the results in terms of residential segregation. For example, Ref. [80] reinterprets Schelling’s model as a model of crystal growth, while Ref. [81] studies scaling properties of the interfaces of the emergent clusters. See also Ref. [82] for an amusing account by Dietrich Stauffer.

2.6. *Scaling theory of cities*

A small city is not a small version of a large city. As mentioned, there has been a considerable hype around *self similarity* and *scale-free* patterns [83] that, to some extent, has fuelled the development of models we have discussed. However, beyond the power-law size distribution, the way a city operates depends on its size. How things depend on size has, for long, been a common theme in biology and ecology. Note the difference to finite-size scaling in statistical physics where the goal is to extrapolate the results to the infinity limit (to study critical phenomena).

Scaling theory has recently come to the attention of physicists [84]. This interest comes from the physical theories of allometric scaling [85]. One of the most influential papers is Ref. [86] that found that different sectors of the economy depend differently on the size of cities. For example, sectors that need people to collaborate—like research and development—scale superlinearly with city size. In contrast, facilities that need to exist relatively close

in space—like gas stations—scale sublinearly. Ref. [87] reestablishes these results in a framework more suitable for physics style modelling by using population density rather than city size as a basis for the scaling analysis. Ref. [88] proposes a model that relates many scaling exponents and finds the regions of parameter space where a city can exist.

2.7. Human mobility

Mobility studies mostly concern how many people move between two locations at a specific time. One can break down this topic in many ways. One can divide the people according to age, sex, or socioeconomic indices. One can separate different times of day, different months, or long-term trends. One can consider moving of residence, work, or the individual themselves.

The origin of human mobility studies is Ernst Ravenstein’s 1885 *The laws of human migration* [89]. Ravenstein noticed, among other things, that the distance people move (their home) follows a skewed distribution—most people move only a short distance.

A more quantitative mobility law is the *gravity law* stating that the number of people T_{ij} travelling between two locations i and j is

$$T_{ij} = c \frac{P_i P_j}{d_{ij}^\delta}, \quad (8)$$

where c and δ are constants, d_{ij} is the distance between i and j , and P_i the population at location i . This relation was first studied by Zipf [90], who only considered the exponent $\delta = 1$. The phrase ‘gravity law’ was coined later in the transportation literature, so it does not appear like Newtonian mechanics played any deeper role in this development than providing a namesake. Subsequent studies have tried to measure and explain the exponent [91] and otherwise improve the gravity law by adding information about the locations [92].

The gravity law was recently improved by the *radiation model* of human travel stating that

$$T_{ij} = T_i \frac{P_i P_j}{(P_i + s_{ij})(P_j + s_{ij})}, \quad (9)$$

where $T_i = \sum_j T_{ij}$ and s_{ij} is the number of people in the circle centred on i and j at the perimeter. The radiation model’s main advantage is that it builds on some simple mechanistic assumptions, whereas the gravity model

is merely a statistical relation. Indeed, the radiation model’s assumptions are rather reminiscent of Stouffer’s theory of intervening opportunities [93], stating:

The number of persons going a given distance is directly proportional to the number of opportunities at that distance and inversely proportional to the number of intervening opportunities.

Stouffer had moving to change work in mind, and the opportunities in question were job opportunities.

2.8. Trajectory analysis

With the advances in position tracking technology over the last couple of decades, researchers have gotten access to large datasets of people’s trajectories. Most commonly, researchers have used datasets from cellphones where people’s locations are identified by the location of the cellphone tower their phone is connected to. From such studies, often comprising hundreds of thousands of individuals, the overarching discovery is just how predictable people are [94, 95, 96]. In most situations of our daily lives, given a sequence of locations visited, one could guess the next location by a probability of around 90% [96]. This phenomenon is also observed in disasters, where people’s routines could be forced to change completely. Still people have been observed to settle into new, highly predictive movement patterns [97].

Another type of trajectory analysis is based on the shape of vehicular travel routes. The most fundamental quantity is the actual travel distance divided by the Euclidean distance between origin and destination. The average value of this quantity, often measured as a function of the Euclidean distance, has many names in the literature, here we follow Ref. [98] and call it *detour index*. For very short travel distances, the detour index could be above two (the travel distance is over twice the straight distance). As the distance increases, the detour index converges after 20-30 km to a value of around 1.3. Many generative models of city maps can reproduce this observation [99].

Another type of study based on the car-travel routes is focusing on how the city shapes the trajectories. Ref. [100], for example, investigates whether the fastest travel routes by cars between points at equal distance from the city centre, tend to move first in, then out, or vice versa. This tendency could be quantified by the *inness*—the area enclosed by the trajectory and




Figure 6: Greedy navigators, agents following a stylised navigation strategy, getting lost in the Leeds Castle maze network. The graph is a ‘visibility graph’ constructed by choosing as few nodes as possible so that the entire study area is visible from at least one node [101]. The edges connect nodes that are visible from one another. The red arrows indicate the actual route travelled by a greedy navigator. The red solid circles denote the shortest path discovered by the greedy navigator.

Source: Reprinted figure from Ref. [102].

the shortest path from origin to destination on the same side the city centre, minus the corresponding area on the other side. Cities dominated by highway ring roads tend to have negative inness because people travel out to the ring road, follow it, then travel in towards the city centre to reach the goal. Ref. [100] measures inness for close to a hundred cities worldwide and relates it to socioeconomic indicators.

2.9. Navigability

Physicists have not only been interested in the structure of actual trajectories in urban car travel, but also how to find the destination when one does not have full information. The way people navigating their surroundings is an active area in cognitive science [103], and it is accepted that some cities, or buildings, are much easier to get lost in than others.

Any attempt to quantify the navigability of a city or building must rest on a model of how people exploit contextual information. Ref. [102], for example, uses a framework called *greedy navigators* in which the individuals have a notion of the direction to their destinations. At every intersection, a greedy navigator chooses the street, not previously travelled, that points

most directly towards the target. In Fig. 6, we show a proof of concept of how greedy navigators fail to find a short path in a garden maze (designed to be hard to navigate). Using greedy navigators one can obtain a navigability index similar to the detour factor—the average distance found by the greedy navigators divided by the actual average shortest distance.

2.10. Future outlook

So far, the physics of urban systems has not been driven by a paramount goal. Instead, it has been building on a collection of observations from data that physicists can, and do, try to explain. However, urban science [48], in a multidisciplinary sense, has some general directions. From an engineering perspective, one would like to make cities sustainable and energy-efficient; whereas seen from social science—because of the ongoing urbanisation of our planet—one would like to foresee the problems and tap into the opportunities of ever-larger metropolises, perhaps via the spatio-socio-semantic analysis framework [104].

3. Traffic flows

Understanding and predicting traffic flow is important for social engineering in general and urban planning in particular [105, 106]. The study of traffic flow in the engineering sciences is thus old, dating back to the 1930s [107]. Outside of the engineering sciences, however, scientists have only recently discovered vehicular traffic as an interesting complex system that perhaps could be described with a few simple laws, and is thus worthy of study with scientific methods [108, 109, 110].

For physicists, traffic-flow models became a topic in the 1990s. In the recent decade, this topic has cooled down somewhat, but nonetheless remains an active field of research in physics and elsewhere, for instance, machine learning [111]. It is probably fair to say that the main motivation for physicists has never been to provide practical advice for urban planners. Rather, the attraction was that vehicular traffic exhibits several types of self-organised, collective phenomena also seen in statistical physics. It is no coincidence that the founding era of the physics studies of traffic flow was in the 1990s. This was a time when there was a prevailing idea that many phenomena in nature and society were connected by underlying ubiquitous organisational principles such as self-organised criticality [83], manifested by many quantities that follow power laws. Even if this view has lately fallen

out of fashion [61], the idea that vehicular traffic is an archetypal complex systems—self-organised, decentralised, and with emergent behaviours that connect short and long spatial and temporal scales—still prevails.

Another reason for studying traffic flows is that the models themselves are interesting. They are among the simplest possible models of emergent phenomena in non-equilibrium systems. Furthermore, they bridge several different modelling frameworks (although none of them originally from physics) [112]—from continuous models of traffic density [113], via discrete particle models (called ‘car-following theories’ in this context) [114], to cellular automata [115].

3.1. Observed phenomena

As mentioned, the primary focus of physicists interested in modelling traffic flow has not been to make accurate forecasting, but rather to qualitatively explain emergent phenomena. So what phenomena can be studied by models? In this section, we will go through some of these observations. Unless stated otherwise, we will discuss phenomena at continuous sections of highways.

Occupancy-flow relations. If the traffic is light, a higher density of cars means that the flow increases—cars move at about the same speed so twice as many cars means twice as large flow. As the traffic gets denser, however, the average speed decreases. Eventually the flow starts decreasing as well. It has been known for over half a century that these quantities do not have a smooth relationship. Since Ref. [116] it is rather thought to be tent-shaped (Fig. 7), or inverse λ -shaped. This suggests the existence of two dynamic phases—a *free-flow* and a *congested* state [108], with an intermediate maximum flow [109].

Sometimes the congested state is divided into *synchronised flow* in which cars are following each other at a relatively constant speed, and *stop-and-go* motion (the name explains the concept) at even higher densities [117]. Some authors go further into dividing the synchronised flow depending on whether the speed and separation of the vehicles is stationary or not [118].

Phantom traffic jams. The distribution of speeds is fairly well described by a Gaussian distribution for all almost all traffic densities (Fig. 8). There could be some anomaly for intermediate speeds (40 km/h in Fig. 8) which could ring a bell for physicists familiar with critical phenomena [119]. Note however that there is nothing scale-free about the speed distribution at this




Figure 7: Observed flow-occupancy relation in highway traffic in Ontario, Canada, 1985. Occupancy is measured (in percentage) as the fraction of time a sensor is blocked by a car. Flow is the number of cars passing per hour. Every data point is an average over 5 minutes.

Source: Reprinted figure from Ref. [116].

point (scale-free, or power-law, distributions are usually the hallmark of phase transitions [61]). Still, there is one supposedly self-organised, emergent, phenomenon believed to explain this anomaly—*phantom traffic jams*. These are jams that happen seemingly without an external trigger. Fig. 9 shows the classical figure illustrating phantom jams with data from aerial photography from Australia in 1967 [120] and reproduced in almost every review paper or book on the subject [108, 109, 121], and also original research papers [122].




Figure 8: The distribution of speeds at different densities. The dashed lines are the best fitting Gaussian distributions.

Source: Reprinted figure from Ref. [109].

Apart from the existence of phantom traffic jams this figure also shows that the jam moves in a direction opposite to the traffic, a finding that has been established by other measurements [123].

Hysteresis. Hysteresis in traffic flow refers to the observation that the average speed-up as traffic gets lighter does not follow the same curve as when traffic gets denser. In the former situation, the average speeds are lower. This was first studied rigorously in Ref. [120] (although mentioned in earlier works).

Pinch effect. There is some evidence that stop-and-go type congestion waves are triggered at special locations along roads where small jams are formed, that later merge to form larger jams (Fig. 10). This is called the *pinch effect* and the larger jams are called *wide moving jams* (although, from a driver's perspective they are rather long than wide).

3.2. Traffic-flow models

Next, we describe several types of traffic-flow models that have been used to explain the observed phenomena. Some of these models have true physics origins, whereas others became popular among physicists, although their origins lie elsewhere (e.g., computer science and mathematics).




Figure 9: Phantom traffic jam in highway traffic. Each line represents one vehicle in a particular lane (trajectories suddenly appearing mean a vehicle made a lane change). Original data from Ref. [120]. Source: Reprinted figure from Ref. [109].

Fluid-dynamical models. *Macroscopic*, or *fluid-dynamic*, models of traffic only use traffic density ρ , flow Q , and average velocity v as variables describing the system. These are related by definition as

$$Q(x, t) = \rho(x, t)v(x, t), \quad (10)$$

where x is the location along the road and t is the time. Assuming continuity (that no cars are generated, or disappearing, along the road) we get the following equation describing mass conservation

$$\frac{\partial \rho}{\partial t} + \frac{\partial Q}{\partial x} = 0. \quad (11)$$




Figure 10: The pinch effect. Panel (a) shows the setup of sensors along a highway. Panel (b) shows the readings (average speeds) of some of these sensors. The pinch effect is manifested in narrow dips (jams) that are created around D5 and aggregate to ‘wide jams’ down the road.

Source: Reprinted figure from Ref. [124].

Eq. (11) should be a part of all fluid-dynamical traffic flow theories, but we need one more equation to make it a full theory. The oldest approach is to assume Q is only a function of ρ and the functional relationship is to be inferred by data

$$Q(x, t) = Q[\rho(x, t)], \quad (12)$$

leading to

$$\frac{\partial \rho}{\partial t} + C(\rho) \frac{\partial \rho}{\partial x} = 0 \quad (13)$$

where $C(\rho)$ comes from data. This *Lighthill-Whitham theory* [113] describes kinematic waves that travel in the opposite direction of the traffic flow (ac-

ording to observations). In the solution of the Lighthill-Whitham equations, shock-waves of infinite density build up. These should be interpreted as jams, and are a challenge numerically, but not conceptually. There are several more sophisticated theories following the footsteps of Lighthill and Whitham. All of them replace Eq. (12) by a more elaborate equation.

Kinetic models. *Kinetic models* of traffic flow are inspired by the kinetic theory of gasses. They use a distribution of car speeds as their main variable describing the state of the system. The original kinetic theory of car traffic was proposed by Prigogine and co-workers in the 1970s [125]. It was quite similar to the original model from physics and to little surprises it shows many discrepancies with car traffic (for example, all cars would drive with the same average speed). This model was later heavily modified by Pavri-Fontana [126]. Like above, Refs. [108, 109] give a summary of these theories.

Car-following theories. We have mentioned traffic flow theories inspired by fluid dynamic and kinetic gas theory, maybe to little surprise, there are also theories inspired by Newtonian mechanics. Such, *car-following theories* are based on equations for the individual drivers and their response to the behaviour of the preceding car. The simplest car-following equation, due to Reuschel [114], is

$$\tau \ddot{x}_n(t) = \dot{x}_{n+1}(t) - \dot{x}_n(t). \quad (14)$$

This equation is derived from assumptions that a driver wants to drive as fast as the preceding car, but not get closer than a safety distance. Later, improved, car-following theories have assumed each car has an internal desired speed, that it follows unless it needs to avoid a collision. These more sophisticated theories can explain the mirrored- λ shape of the flow-density curves and hysteresis effects.

Coupled-map lattice models. The models we have seen so far have all been continuous in both time and space. So called *coupled-map lattice models* share many assumptions of car-following theories, but use a discrete time. In general, such models have the form

$$v_n(t+1) = \text{Map}_n(v_n(t), v_{n,\text{des}}, \Delta x_n) \quad (15a)$$

$$x_n(t+1) = v_n(t) + x_n(t), \quad (15b)$$

where $\text{Map}_n(\cdot)$ is a dynamical map that takes into account the speed and position of the n th vehicle, v_n and x_n , the desired speed of the n th vehicle,

$v_{n,\text{des}}$, and the headway to the $n+1$ th vehicle, Δx_n . This versatile framework can accommodate different personalities of drivers, and different classes of vehicles, etc. Several of the empirical characteristics of traffic flow (such as flow-density relations) can be reproduced by coupled-map lattice models. Popular models of this kind includes those of Yukawa and Kikuchi [127] and Krauss, Wagner, and Gawron [128].

Cellular automata models. One further abstraction from coupled-map lattice models is to discretise space as well as time. This leads to so called cellular automata models. These are the most well-studied type of models in the physics literature, which might seem surprising since it is a type of models derived from computer science and mathematics rather than physics (whereas the kinetic and car-following models above have a much stronger physics flavour). The explanation is probably that physicists became interested in traffic models as a part of a general hype around complex systems, where cellular automata models of artificial life are among the most iconic theories.

The most well-studied cellular automata model of traffic flow is the Nagel-Schreckenberg model [115]. In this model, the road is represented by a one-dimensional discrete lattice. There are N vehicles on this road. Each cell of the road is occupied by maximally one vehicle. All vehicles are updated in parallel according to the following rules (to be followed in order):

1. *Acceleration.* If $v_n < v_{\text{max}}$, then the speed of vehicle n is increased by one unit, otherwise the speed is unchanged:

$$v_n(t+1) = \min(v_n(t) + 1, v_{\text{max}}). \quad (16)$$

2. *Deceleration.* If $x_{n+1} \leq x_n + v_n$ —that is, the car ahead is so close that vehicle n would reach its position (or further) the next time step—then the n th vehicle brakes:

$$v_n(t+1) = \min(v_n(t), x_{n+1} - x_n - 1). \quad (17)$$

3. *Randomisation.* By chance, that is, with probability p , the speed of some cars is decreased:

$$v_n(t+1) = \max(v_n(t) - 1, 0). \quad (18)$$

4. *Vehicle movement.* Each vehicle moves forward according to its new speed:

$$x_n(t+1) = x_n(t) + v_n(t+1). \quad (19)$$

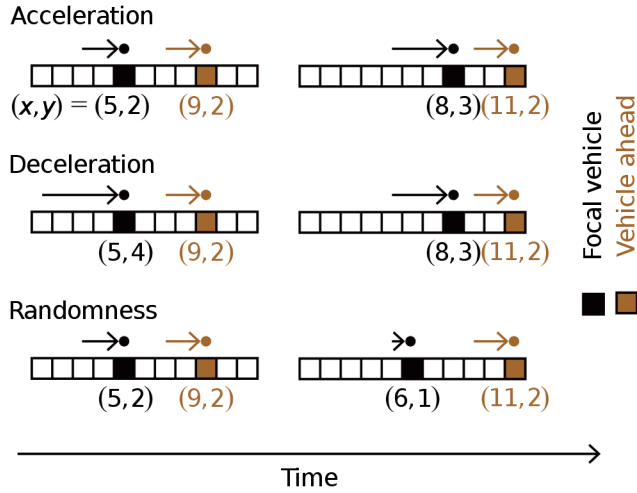


Figure 11: The three rules for changing the speed in the Nagel-Schreckenberg model. The rules are illustrated for the focal vehicle, while the vehicle ahead keeps its speed. The arrow illustrates the speed v ; the circle illustrates the position x .

See Fig. 11 for an illustration of the Nagel-Schreckenberg rules.

The Nagel-Schreckenberg model can, despite its simplicity, reproduce many features of real traffic, such as the flow-density curves and phantom traffic jams. With this model as a starting point the research has branched out in many directions. Some of the research has striven to include more realism [129, 130], while other [131] has shown that it takes only a small modification to turn it to a model of self-organized criticality (the Nagel-Schreckenberg model itself does not have the necessary meta-stable state). Yet others studied further simplified models as discussed below, although these simplified models are incapable of reproducing the above-mentioned full statistical characteristics of traffic.

Connections to non-linear statistical mechanics. If one gives up on trying to reproduce all statistical features of highway traffic, then one can further simplify models like the Nagel-Schreckenberg cellular automaton. This will typically reduce the models to standard models of non-equilibrium statistical mechanics, like the *totally asymmetric simple exclusion process* (TASEP) [132] or the Burgers' equation [133] (in particular, its noisy version [134]). These more stylised studies are often focused on finding dynamical critical exponents that relate the size of a system to its dynamics and the critical point

separating the free flow and congested phases (see Ref. [135] for a typical example).

3.3. *Pedestrian traffic flows*

Pedestrian traffic is a related, but far from an equivalent type of socio-physical system compared to vehicular flow. It could also be thought of as a self-organised granular flow of semi-intelligent particles. The main research questions concern the formation of trails in open landscapes [136], the formation of lanes in dense pedestrian traffic [137], and escape panic behaviour [138].

Most models of pedestrian flows take their inspiration in physics and model the individuals as particles driven by forces [139]. Several things are different from real gasses—there is, for example, no conservation of momentum. Typically one has to assume that people repel each other by two forces; one is social—the desire not to be too close to another person—and one is physical—crowded conditions such that people actually have to be in physical contact [109]. To accurately model escape panic, one has to break down the physical forces into tangential and radial components.

3.4. *Future outlook*

From a physics point of view, the field of traffic flow modelling seems to have somewhat cooled down at the time of writing. Elsewhere, it is a topic of emergent interest. In particular, the recent boom in research on autonomous vehicles has renewed the interest in applying machine learning to these topics [111]. The current interest in self-driving cars produces an enormous amount of data. Most of it arguably useless for this type of research, but probably eventually enough to discover new statistical laws of vehicular traffic. Even if data does not come as a side product from the automotive industry, it is nowadays easier and cheaper to collect. There have been projects to this end that rely on drones [140] or tower-mounted cameras [141]. In Fig. 12, we plot individual trajectories of some of the recordings from one of these datasets [140].

When we—by new observations from new datasets—have created new qualitative statistical laws to replace the current qualitative observations, then the question will once again be to find minimal models recreating the observations. In particular, with high-quality data on the onset of the congested phase, we could measure how often phantom traffic jams actually




Figure 12: Trajectories, from the highD dataset [140] plotted in the style of Fig. 9. **A**, Free-flow state (recording 60, lane 8). **B**, **C**, More congested traffic with stop-and-go waves (recordings 26 and 25, lanes 2 and 4, respectively).

occur and whether the current mechanistic models of these can explain the observation, or if we have to go back to the drawing board.

4. Econophysics

Unlike traditional economics, which is built upon a rational-choice model, econophysics borrows the particle model from statistical physics to explain the behaviour of an agent. Such a model assumes that the agent's tastes and preferences are not fixed, but instead depend on the interactions with other agents [142]. In other words, econophysics puts a greater emphasis on the social environment of the agent [143]. Some other physics models and concepts commonly applied to economics include the kinetic theory of gases, chaos theory, percolations, and self-organised criticality.

Empirical work in econophysics is mostly focused on the analysis of firm growth and competition, industry entry and exit rates, money flows, financial markets, and international trade [144, 145, 146, 147, 148, 149, 150], that is, on areas in which huge datasets are available, and the application of statistical-physics tools and methods proves useful. The areas of economics with scarce data availability, such as macroeconomics in which datasets are short and noisy, has not attracted much attention among econophysicists. However, with the increasing acceptance of networks in the mainstream economics, econophysics may still play an important role in the future development of macroeconomics [143].

4.1. *The advent of econophysics*

In 1991, Mantegna published a paper in a physics journal, *Physica A* [151], in which a time series of daily financial market prices was analysed and price changes were shown to follow a power-law distribution. At that time, power-law distributions and scaling relations were attractive key topics to statistical physicists after the big booms of *fractals* in the early 1980s [152, 153, 154] and *self-organised criticality* in the late 1980s [155, 83]. Research targets of interest to physicists were extended widely to general complexity in nature, thus crossing the traditional boundaries between research fields. Market price changes were a part of this extension and got accepted as one of physically interesting phenomena that exhibit power-law behaviour.

The next pioneering interdisciplinary paper appeared in the same journal in 1992 by H. Takayasu et al. [156]. A simple artificial model of the market was proposed comprising mathematically defined dealers in the form of dynamical particles in a one-dimensional space of prices. The model's non-linear dynamics caused chaotic time evolution resulting in almost random price movements (see the next subsection for details).

In 1995, Mantegna and Stanley analysed the time series of stock-market prices recorded at a one-minute interval, and found that the price-change distributions at different time scales, upon re-scaling, conform to a function with symmetric power-law tails [157]. Such a data-analysis method was familiar from the study of critical phenomena involving phase transitions. Because their paper was published in a high-impact journal, the new physics approach to financial markets attracted wide attention. In the same year, Stanley coined the term *econophysics* to represent an interdisciplinary research field that focuses on economic phenomena from a physics point of view. He introduced this term at a conference on statistical physics held in Kolkata, India.

In 1997, the first international meeting with the title ‘econophysics’ was held in Budapest, Hungary. Most of the gathered researchers were specialists in statistical physics, although there was notable participation from fields as diverse as high-energy experiments. In this same year, the journal *Physical Review Letters* also opened the door to econophysics, and a theoretical paper explaining the generating mechanism for power-law distributions in financial markets was published [158]. Prior to this acceptance, the journal was rejecting econophysics manuscripts on the basis of the topic being out of scope. The acceptance thus marked the promotion of econophysics to a status of a fully fledged field in applied physics, such as biophysics or geophysics. The number of econophysics researchers subsequently increased, as did the variety of research topics that began to cover more than just financial markets.

In 1999, the first monograph on econophysics was published [159], and textbooks on financial markets from the physics viewpoint followed [160, 161]. Multiple workshops and conferences on econophysics were held annually thereafter, with many economists and finance researchers joining to discuss practical problems [162, 163, 164].

The rest of this chapter focuses on the development of econophysics research stemming from the afore-mentioned simple physical model of financial markets [156]. We first describe the historical background of the dealer model, and show how more advanced dealer models have arisen. Then we introduce an empirically derived time-series model called the PUCK model, and proceed to demonstrate the relation between multiple dealer models. We also outline a recent analysis of comprehensive market data that includes all microscopic orders appearing on the Foreign Exchange market at a millisecond interval. The mechanism of financial Brownian motion is compared with the physical phenomenon of colloidal Brownian motion, and the most ad-

vanced dealer model, which is reconstructed directly from the data, is solved by applying the classical kinetic theory. We close off the chapter with a discussion of a novel emerging perspective, that of an ecosystem of strategic dealers.

4.2. Agent-based modelling: The dealer model

Mandelbrot's inspiration for introducing the concept of fractals, that is, the scale invariance of complicated shapes in nature, originated during an examination of historical cotton-price charts at various time scales. Market prices have thus become the very first example of fractals. In part through their interactions with Mandelbrot, H. Takayasu and Hamada—a physicist and an economist—joined forces to create a model of financial markets that would explain why market prices fluctuate in a scale-invariant manner [156, 165, 166]. A prevailing view in economics at the time was that if all market dealers were rational and possessed enough information, then the market price would be uniquely determined and stable. Yet, this view could not be further from the real-world price fluctuations, which prompted H. Takayasu and Hamada to construct their model borrowing ideas from physics. The model thus had to incorporate the essential underlying mechanisms and processes in a way that is as simple as possible, but non-trivial.

Let us envision an artificial financial market comprising N dealers who buy and sell financial instruments such as stocks. Every dealer is assumed to try to buy at a low price and sell at a high price, hoping to earn the price difference. The i th dealer's trading action at time t is described by introducing two threshold prices, the buying price, $b_i(t)$, and the selling price, $s_i(t)$. The dealer hopes to buy at the former price or lower, and to sell at the latter price or higher. The difference, $L_i(t) = s_i(t) - b_i(t) > 0$, called the spread characterises greediness of the dealer, which is set to a constant value L in the simplest case. All dealers' buying and selling prices are gathered to make the market's order-book. A deal occurs if the condition $s_i(t) \leq b_j(t)$ is fulfilled for a pair of dealers i and j , in which case the i th dealer sells to the j th dealer, or equivalently, the j th dealer buys from the i th dealer. An interesting point is that no deals occur if all dealers' buying prices are within the distance L from the minimum buying price. Only when the distance between the farthest pair of dealers equals or exceeds L can a deal (between this particular pair) take place. Deals thus represent a strong, non-linear, attractive interaction that makes a group of dealers compact in the price space.

The model is further simplified by assuming that dealers can possess at most one stock at a time. If a dealer possesses a stock, they are a seller quoting only the selling price. If the dealer does not possess a stock, then they are a buyer quoting only the buying price. A seller hopes to sell at a high price, but if there is no buyer whose buying price is equal or higher, then the seller should compromise by lowering the price until a trade becomes possible. This situation is described by a differential equation

$$\frac{db_i(t)}{dt} = \sigma_i(t)c_i, \quad (20)$$

where $\sigma_i(t) = 1$ ($\sigma_i(t) = -1$) signifies the buyer (seller) state of the i th dealer, while $c_i > 0$ quantifies the dealer's (initially random) hastiness. A deal between the seller i and the buyer j occurs when $s_i(t) = b_i(t) + L \leq b_j(t)$, at which moment the state functions $\sigma_i(t)$ and $\sigma_j(t)$ change their signs. The resulting market price, $P(t)$, which takes the value of the latest deal, evolves deterministically in time. The model is initialised such that all dealers start with the same buying price, and some dealers start as buyers and others as sellers.

With $N = 2$ dealers, one seller and one buyer, the time evolution of the market price is almost trivial. The two dealers periodically alternate their state and the resulting market price oscillates regularly. With $N \geq 3$ dealers, the time evolution of the market price becomes highly non-linear. There is, in fact, an underlying chaotic effect that magnifies small initial differences. We thus learn from this simple model that even fully deterministic dealer behaviour can cause noisy dynamics. The model is, nonetheless, insufficient to explain the fractal properties of market prices.

A minimal modification of the described model accounts for an effect called trend following. A dealer who follows the trend expects that price movements keep moving in the same direction as in the immediate past. This can be mathematically formulated using a moving average of length T

$$\langle \Delta P(t) \rangle = \int_0^T w(u) \frac{d}{dt} P(t-u) du, \quad (21)$$

where $w(u)$ is a weight function that satisfies $\int_0^T w(u) du = 1$. Eq. (20) is then appended with a trend-following term

$$\frac{db_i(t)}{dt} = \sigma_i(t)c_i + d_i \langle \Delta P(t) \rangle. \quad (22)$$

Coefficients d_i quantify the extent of trend following. They are usually positive for dealers who are trend followers, but can also be negative for dealers called contrarians. The simplest possible variant of the model is obtained by assuming that all dealers are trend followers with the same coefficient $d > 0$. Despite being simplistic, this assumption has a drastic effect, yielding deterministic price dynamics that resemble scale-invariant fluctuations of stochastic random walks [156]. The model thus identifies two mechanisms likely to be responsible for some of the key characteristics of realistic market-price time series. Namely, market prices fluctuate almost randomly due to non-linear, chaos-inducing interactions between dealers, while scale-invariance emerges from the dealer tendency to follow trends.

The dealer model with trend following as defined by Eq. (22) can be studied analytically to some extent. The dynamics of the centre of dealer mass follows a Langevin equation [165], which is an equation that is well-known in the context of colloidal Brownian motion. The distribution of market-price changes obeys a power law with an exponent that depends on the value of the parameter d [167], in line with known empirical facts [151, 159] and a previous theoretical analysis [158]. The dealer model can also be used as a basis for deriving the autoregressive conditional heteroskedasticity (ARCH) model of Engle [168], thus offering a mechanistic explanation for the origins of volatility clustering observed in financial markets [169]. Finally, by tuning parameters values, the model reproduces the phenomenon of abnormal diffusion, as well as the statistical properties of dealing time intervals [166].

A stochastic variant of the dealer model was introduced in 2009 [170]. In the model, the function $\sigma_i(t)$ is random, such that—at each moment of time t —either $\sigma_i(t) = 1$ or $\sigma_i(t) = -1$ with the probability of 0.5. This modification improves upon already favourable properties of the dealer model even in the case of $N = 2$, which is easily solvable using both analytical and numerical methods. The stochastic dealer model can, for example, generate bubble-like behaviours that cause the market price to grow exponentially if the trend-following coefficient, d , is above a certain value.

Further generalisation of the stochastic dealer model has enabled capturing the characteristics of an intervention by the Bank of Japan in the foreign exchange market between the US dollar and the Japanese yen [171]. Aside from ordinary dealers responsible for usual market-price fluctuations, the model also includes a special dealer that takes the role of the Bank of Japan. The special dealer can cause large market-price changes that, according to empirical analyses of market data, are accompanied by risk-averse responses

such as bid-ask spread widening, loss cutting, and profit booking. Ordinary dealers in the model, with some adjustment, can mimic risk-averse responses and thus generate said empirical phenomena in simulations. The increased realism makes it possible to assimilate financial time-series data into the model. This opens the door to the planning of intervention strategies, as well as predicting subsequent market responses.

More recent extensions of the stochastic dealer model help clarify the cross-currency correlations between the US dollar, the Japanese yen, and the Euro [172]. A new type of dealers, who pursue what is known as triangular arbitrage, is introduced into the model. Such dealers earn profit by quick circular exchange transactions from, for example, the US dollar to the Japanese yen to the Euro and back to the US dollar. Interestingly, triangular arbitrage in currency markets was first reported in 2002 in an econophysics study [173]. New evidence shows that such arbitrage still exists in the present financial markets in which automated trading systems dominate [174]. The stochastic dealer model has clarified that it is, in fact, a small number of triangular-arbitrage dealers who boost the cross-currency correlations in a manner consistent with empirical observations.

4.3. Time-series modelling: The PUCK model

Trend following introduced in the dealer model is based on the idea that dealers make their decision referring to the latest market data using the moving average. Applying a similar idea to the time series of deal intervals, known for their temporal-clustering behaviour [175], it was found that the occurrence of deals in markets is modelled well by a Poisson process with a time-dependent mean value. This value is given by the moving average of the latest deal intervals over a time period T , where the best estimate of $T = 150$ s was obtained from the dollar-yen exchange-market data at the time. The finding clarifies the mechanism underlying the temporal clustering of deal intervals. When random fluctuations cause a few short intervals to repeat, the moving average value becomes smaller, making shorter intervals more likely to appear in the Poisson process. A dense period with short deal intervals ensues. Converse is true when random fluctuations cause a few long intervals to repeat. The described time-dependent Poisson process is a self-modulation process whose fluctuations have a $1/f$ power spectrum bordering between stationary and non-stationary processes [176].

The idea of using the moving average was also extended to the time-series analysis of market prices [177, 178]. For a given time series of real market

prices with a fixed sampling interval, $\{p(t)\}$, the following time-evolution model is defined

$$p(t+1) - p(t) = -\frac{b(t)}{M-1}[p(t) - p_M(t)] + g(t), \quad (23)$$

where M is the length of the moving average, $p_M(t) = \frac{1}{M} \sum_{k=0}^{M-1} p(t-k)$, and $g(t)$ is independent noise with a zero mean. The coefficient $b(t)$ denotes a slowly changing parameter estimated from the time-series data. The case of $b(t) = 0$ corresponds to the ordinary random walk. In the case of $b(t) > 0$, the future market price, $p(t+1)$, is likely to be attracted to the latest moving average price, $p_M(t)$, signifying stable market-price movements. In the case of $b(t) < 0$, the future market price is likely to be repelled from the moving-average price, signifying unstable market-price movements.

The described model can be generalised with a time-dependent market-potential function, $\Phi_M(t)$, as follows

$$p(t+1) - p(t) = -\left. \frac{\partial}{\partial p} \Phi_M(t) \right|_{p=\frac{p(t)-p_M(t)}{M-1}} + g(t). \quad (24)$$

This is the Potential of Unbalanced Complex Kinetics (PUCK) model. Eq. (23) is a special case of the PUCK model with the quadratic potential $\Phi_M(t) = \frac{1}{2}b(t)p^2$. The PUCK model with the quadratic potential function has been shown to apply to various financial markets, successfully reproducing most of empirical stylised facts such as the power-law distribution of price changes, abnormal diffusion over short time scales, as well as volatility clustering [179].

The market-potential function is estimable for any market-price time series, including those artificially produced by the dealer model described in Section 4.2. A stable quadratic potential is obtained for contrarian dealers ($d < 0$), an unstable quadratic potential is obtained for trend followers ($d > 0$), and an asymmetric higher-order potential appears when trend following is asymmetric (Fig. 13). Moreover, the value of the market-potential coefficient $b(t)$ can be theoretically derived from the dealer model, thus demonstrating that the origin of the market-potential function in Eq. (24) comes from the trend-following behaviour of dealers.

A merit of the PUCK model is its wide applicability; the model describes market-price time series in various circumstances. This ranges from nearly random walks under normal market conditions, an exponential divergence in the case of bubbles or crashes, and even a double exponential divergence




Figure 13: An example of artificial market-price fluctuations produced by the stochastic dealer model with the corresponding estimates of the potential functions. The parameter $d = d_1 = d_2$ from Eq. (22) varies such that $d = -1$ during the first 1000 time steps, $d = 0$ during the second 1000 time steps, $d = 1$ during the third 1000 time steps, and $d = -1$ if $\langle \Delta P \rangle_M < 0$ whereas $d = 1$ if $\langle \Delta P \rangle_M \geq 0$ during the last 1000 time steps, where $\langle \Delta P \rangle_M$ is the moving average from Eq. (21) of length $M = 10$ ticks with uniform weights. Source: Reprinted figure from Ref. [170].

or a finite-time singularity in the case of hyper inflation [180, 181]. The threshold between the normal and the abnormal random walk is the value $b(t) = -2$ of the quadratic-potential coefficient; namely, when $b(t) < -2$, Eq. (23) becomes linearly unstable causing price fluctuations to grow or decline exponentially as is observed in bubbles or crashes, respectively. If a cubic potential function is detected, it generally corresponds to asymmetric price movements [182].

In a short time-scale limit, Eq. (23) reduces to the Langevin equation for Brownian motion containing a mass term and a viscosity term [180]. Interestingly, the mass term is proportional to $-b(t)$, showing that trend

following works as inertia. Also, the viscosity term becomes negative when $b(t) < -2$, suggesting that bubbles, crashes, and inflation should be regarded as negative viscosity phenomena, that is, as being under the influence of an accelerating instead of a decelerating force.

4.4. Order-book modelling: *Financial Brownian motion*

Another approach that physicists brought to financial markets is data analysis and modelling of an order book. Such a book lists buy orders (bids) and sell orders (asks) gathered in a market. Ref. [183] introduced a theoretical model in which bids and asks are injected randomly onto a price axis. A deal occurs when a new bid price is equal to or higher than the lowest ask price, or vice versa, when a new ask price is equal to or lower than the highest bid price. The deal signifies that the corresponding pair of orders annihilate forming the latest market price. Otherwise, injected orders accumulate on the price axis of the order book. The mechanism was pointed out to be similar to a one-dimensional catalytic chemical reaction in which the reaction front moves randomly. Ref. [184] described the details of a stock-market order book, documenting empirical statistical laws about the injection of bids and asks, as well as cancel orders. A simple mathematical model called Zero Intelligence was proposed to capture the empirical findings, with the name of the model stemming from the fact that no intelligent dealer strategy was needed.

Ref. [185] introduced a novel data analysis of order books, focusing on an analogy between colloidal random walks and financial market-price movements. Accumulated bid and ask orders are regarded as water molecules in this analysis, while an imaginary colloidal particle is assumed to exist in the gap between bids and asks centred right in the middle between the highest bid and the lowest ask (Fig. 14). This colloidal-particle picture is intuitive in the following sense. As the particle gets displaced, say, to the right (i.e., towards higher prices), the opened up space in its wake gets quickly filled with water molecules (i.e., bids) from further back where the number of molecules decreases (Fig. 14B, C). In front of the colloidal particle, by contrast, water molecules get pushed forward, decreasing their number next to the particle, but increasing the number further away (Fig. 14B, C). This intuitive picture is fully consistent with the dealer model and trend following by which pairs of buy and sell orders move together with the market price.

A more conventional picture treats all buy orders (and separately all sell orders) on an equal footing, but this is incorrect. As the colloidal-particle

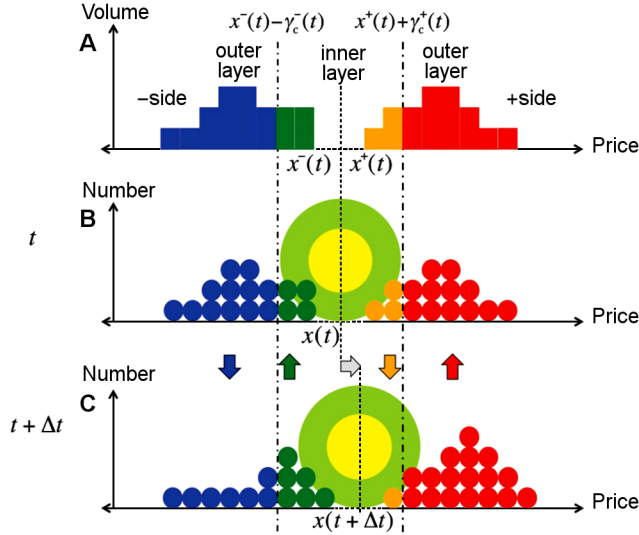


Figure 14: Schematic representation of the Financial Brownian Particle model. **A**, An order-book configuration of buy orders (blue in the outer layer and green in the inner layer) and sell orders (red in the outer layer and orange in the inner layer) on the price axis. **B**, Corresponding configuration of outer-layer particles (blue and red disks), inner-layer particles (green and orange disks), the colloidal Brownian particle’s interaction range (green ring), and the core (yellow circle). **C**, After time Δt the configuration of surrounding particles changes.

Source: Reprinted figure from Ref. [185] under the Creative Commons Attribution 3.0 Unported (CC BY 3.0).

picture shows, buy and sell orders should be categorised into an inner and an outer layer of opposite behaviour. If bids (or asks) increase in the inner layer they decrease in the outer layer and vice versa. Inner-layer orders furthermore play a role of a driving force behind market-price movements, as indicated by a high positive cross-correlation between the velocity of price movements and the rate of change of orders in the inner layer (Fig. 15). Interestingly, outer-layer orders exhibit a negative cross-correlation between the velocity of price movements and the rate of change of orders, and thus can be considered as drag resistance for market-price movements (Fig. 15). All this implies a fluctuation-dissipation relation for the colloidal particle, which is modelled by the Langevin equation. The value of the drag coefficient normalised by the colloidal-particle mass can then be estimated from market-price data [185].

An often overlooked aspect of modelling financial markets using continuous-price models, such as the Langevin equation, is whether the continuous-price

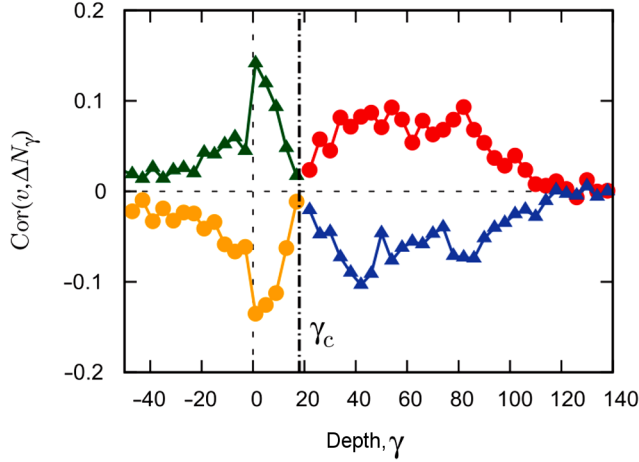


Figure 15: Cross-correlation function between the velocity of market-price movements and the rate of change of orders as a function of depth. Buy orders are indicated with green and blue triangles, whereas sell orders are indicated with orange and red circles. At $\gamma_c = 18$ the nature of the cross-correlation changes, thus revealing the distinction between inner and outer layers. A positive inner-layer (outer-layer) cross-correlation means that orders in this layer act as a driving force (drag resistance) for market movements. Source: Reprinted figure from Ref. [185] under the Creative Commons Attribution 3.0 Unported (CC BY 3.0).

assumption can be justified. To this end, Ref. [186] uses the described analogy between a market-order book and a molecular fluid to estimate the financial Knudsen number. Because the Knudsen number is generally defined as a ratio of the mean free path to a representative length scale of the system, in the case of financial markets, the former is given by the average distance of price movements in one direction, while the latter is the diameter of the colloidal particle in terms of the inner-layer width for both buy and sell sides (Fig. 14). The continuous-price assumption is valid if the Knudsen number is smaller than 0.01, whereas a discrete-time description is needed if the Knudsen number is larger than 0.1 (with transitional regimes in between). The estimated value of the Knudsen number for dollar-yen and dollar-Euro markets fluctuates around 0.05 most of the time, becoming larger than 0.1 in the times of market turmoil. This result indicates that the continuous-price assumption is questionable for modelling financial markets, especially when large market-price fluctuations take place.

4.5. A kinetic approach to financial market microstructure

We have reviewed the dealer model as a financial microscopic model describing decision-making process on the level of individual agents. This model has the advantage that (i) it can capture the strategic decision-making process of individual dealers (such as trend following), and (ii) it reduces to the PUCK model as its macroscopic dynamics for $N = 2$ and thus can replicate the empirical facts seen in the price time series. The dealer microscopic model, however, has a disadvantage that it could not be directly validated from data, because it requires the truly microscopic data to track all traders' decision-making dynamics. Indeed, the trend-following mechanism is theoretically assumed in the model as a strategy of individual traders, which could not be directly confirmed. Also, this model requires calibration of the buy-sell spread distribution. This situation is in contrast to other mesoscopic models, such as purely-random order-book models [183, 184, 187, 188] (see Refs. [189, 190] for reviews), which require fewer calibration parameters although they cannot capture the strategic decision-making process of individual traders.

Recently, the situation with respect to data availability has drastically changed; truly microscopic data has become available due to the big-data revolution, which has enabled confirming various theoretical assumptions of the dealer model, such as the trend-following mechanism and the buy-sell spread distributions. In the following, we review several results of the microscopic empirical analyses as performed in Refs. [191, 192, 193] that summarise the trading strategies employed by real high-frequency traders. In particular, we focus on the trend-following strategies implemented by market makers that directly validate the dealer model with the microscopic data.

Microscopic data: trading logs of individual traders. Here we describe the microscopic data used in the analyses in Refs. [191, 192, 193]. The trading-log data originates from the Electronic Broking Services market, one of the biggest foreign exchange markets in the world managed by the CME Group. This data includes the decision-making process of traders, such as order submissions, cancellations, and executions, with anonymised trader identifiers and anonymised bank codes. Our focus, in particular, is on the exchange market between the US dollar (USD) and the Japanese yen (JPY) from 18:00 GMT on 5 June to 22:00 GMT on 10 June 2016, with the minimum volume unit being \$1MUSD, the minimum price precision (called tick size)




Figure 16: Sample trajectories of the three top HFTs. Plotted are the limit order prices for both bid (blue) and ask (red) sides. The insets illustrate the probability density functions of the buy-sell spread \hat{L}_i for each HFT individually.

Source: Reprinted figure from Ref. [191] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

¥0.005 JPY, and the minimum time precision 1 ms. For brevity, ¥0.001 JPY is used as a price unit called tenth pip or simply tpiip.

Our attention is centred around high-frequency traders (HFTs), typically machines that frequently submit and cancel their orders according to some strategic algorithm. HFTs are defined according to the total number of limit-order submissions. Specifically, a trader who submits more than 2,500 orders weekly qualifies as an HFT. This definition is similar to the one from a previous study [194] of the Electronic Broking Service market. There are many potential alternative definitions that could be considered, but ours offers clarity and the ease of implementation. With this definition in mind, we identified 134 HFTs during the week under consideration. The total number of traders was 1,015.

HFTs as liquidity providers. By plotting three sample trajectories of HFTs in terms of their limit orders (Fig. 16), we observe that the HFTs typically maintain two-sided (buy-low and sell-high) quotes. Generally, two-sided quotes attempt to profit from the bid-ask spread, but are also subject to liquidity rebates that may exceed the trading fees, allowing HFTs to trade with zero marginal cost [195]. In our case, the HFT behaviour is indeed interpreted as liquidity provision (i.e., market making) in response to the request by the platform managers. HFTs have an incentive to play the role of liquidity providers according to the rule book of the Electronic Broking Services market [196].

Here, we denote the best bid and ask prices of the i th HFT by \hat{b}_i and

\hat{a}_i , respectively, where the index i is allocated according to the number of submissions during the week. Any variable with a hat (e.g., \hat{A}) implies a stochastic variable, to distinguish from a real number (such as A). The difference between the best bid and ask prices $\hat{L}_i = \hat{b}_i - \hat{a}_i$ is called the buy-sell (sometimes also bid-ask) spread of the i th HFT. The buy-sell spread \hat{L}_i can be directly measured in our dataset at the level of individual HFTs (see the insets in Fig. 16). In addition, we can define the mid-price of the i th HFT as $\hat{z}_i = \frac{1}{2}(\hat{b}_i + \hat{a}_i)$. The mid-price \hat{z}_i can be interpreted as the appropriate price in the eyes of the i th HFT at the time, while \hat{L}_i can be interpreted as a profit estimate for a round-trip trade, or alternatively, a risk evaluation against adverse selection (i.e., a possibility that the HFT misses some information).

Trend-following behaviour. As discussed in Section 4.2, the dealer model was originally constructed with the assumption of trend-following behaviour of traders in mind. We have validated this theoretical assumption by direct observation, that is, by identifying a statistical correlation in microscopic data that can be interpreted as trend following at the level of individual HFTs. The analytical framework in this context can be described as follows.

First, let us introduce the tick time T as an integer time incremented by 1 when a transaction is executed (see Fig. 17A). We note that the tick time can be mapped onto the physical time as $t = \hat{t}[T]$, where the square bracket stresses that the argument of the stochastic variable is based on the tick time. The mid-price of the i th HFT at the tick time T is represented by $\hat{z}_i[T]$ and the market transaction price is represented by $\hat{p}[T]$.

Then, let us study the correlation between the one-tick future change of mid-price for the i th HFT, $\Delta\hat{z}_i[T] = \hat{z}_i[T+1] - \hat{z}_i[T]$, and the one-tick historical market-price change, $\Delta\hat{p}[T-1] = \hat{p}[T] - \hat{p}[T-1]$ (see Fig. 17A). The average of $\Delta\hat{z}_i[T]$ conditional on $\Delta\hat{p}[T-1]$, denoted $\langle\Delta\hat{z}_i\rangle_{\Delta p}$, for two sample HFTs (Fig. 17B, left panel) shows that the correlation is linear for $\Delta p \rightarrow 0$, but saturates for $\Delta p \rightarrow \infty$. This suggests, on average, a hyperbolic-tangent scaling relation

$$\langle\Delta\hat{z}_i[T]\rangle_{\Delta p} \approx c_i \tanh \frac{\Delta p}{\Delta p_i^*} \quad (25)$$

with the parameters c_i and Δp_i^* unique to the i th HFT. Here, the bracket $\langle\hat{A}\rangle_{\Delta p} = \langle\hat{A}\rangle_{\Delta\hat{p}[T-1]=\Delta p}$ implies the ensemble average of \hat{A} conditional on the previous price change begin fixed to $\Delta\hat{p}[T-1] = \Delta p$ and on the HFT being active, that is, $\Delta\hat{z}_i[T] \neq 0$. Indeed, by re-scaling the horizontal axis to

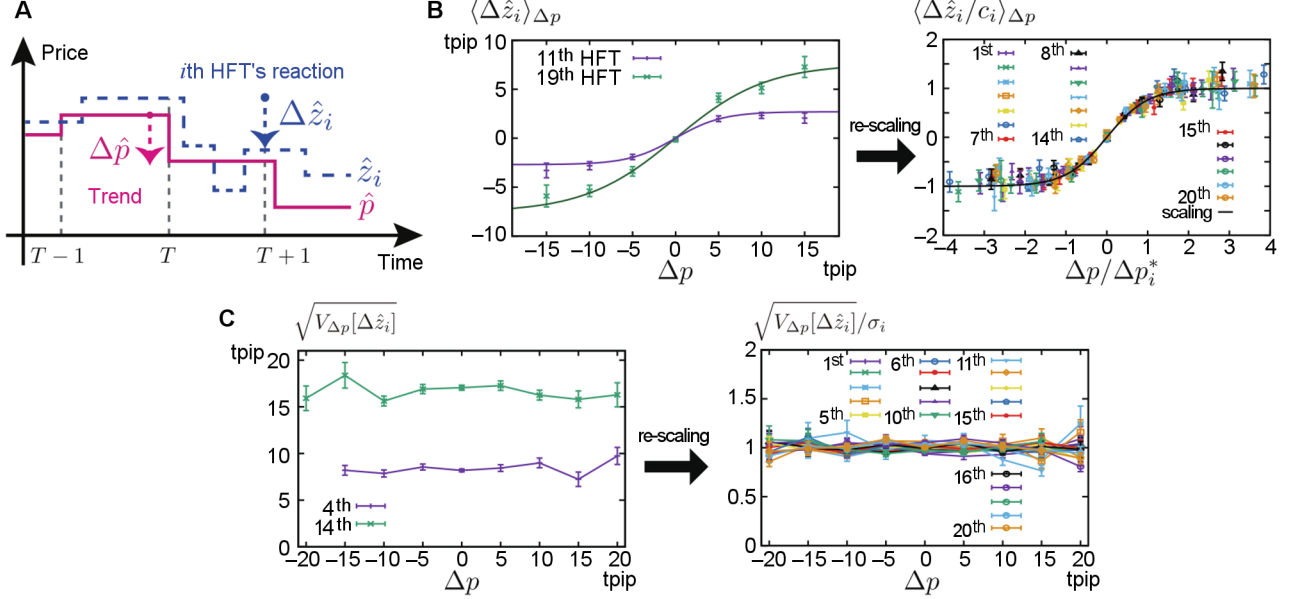


Figure 17: Trend-following analysis. **A**, Tick time T is an integer time measure incremented by 1 at every transaction. $\Delta \hat{z}_i[T]$ is the future mid-price change of the i th HFT and $\Delta \hat{p}[T - 1]$ is the historical market-price change. **B**, Statistical correlation suggests trend-following behaviour on average at the level of individual HFTs (left panel emphasises the 11th and the 19th HFT). Upon re-scaling, the same master curve Eq. (25) is seen to be valid for at least the top 20 HFTs (right panel). **C**, Variance $V_{\Delta p}[\Delta \hat{z}_i]$ conditional on the historical price change before and after scaling (left and right panel, respectively), implying that the variance is irrelevant to the historical price change. Source: Reprinted figure from Ref. [191] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

$\Delta p / \Delta p_i^*$ and the vertical axis to $\langle \Delta \hat{z}_i / c_i \rangle_{\Delta p}$, we observe a clear master curve among the top 20 HFTs (Fig. 17B, right panel), suggesting the universal validity of the formula (25) for the top HFTs in the studied market.

There is also evidence of another scaling relation that holds for the variance of an HFT's one-tick future change of mid-price conditional on the historical market-price change being Δp (Fig. 17C). We specifically have

$$V_{\Delta p}[\Delta \hat{z}_i[T]] = \langle (\Delta \hat{z}_i[T] - \langle \hat{z}_i[T] \rangle_{\Delta p})^2 \rangle_{\Delta p} \approx \sigma_i^2, \quad (26)$$

where the quantity σ_i^2 is a constant unique to the i th HFT. This relation suggests that the variance, unlike the mean, is independent of the historical market-price change Δp ; HFTs follow the trend, but how much they adjust

their mid-price in doing so is solely their intrinsic property.

The scaling relations described here are statistical laws that reveal strategic trading behaviour beyond the previously mentioned zero-intelligence models. Such behaviour holds for HFTs as market makers, but does it differ from what low-frequency traders (LFTs) do?

Indeed, there are noticeable differences between HFTs and LFTs. The former keep a few live orders, typically less than 10, while allocating one unit of volume per order (see Fig. 18A, B). This volume per order is in contrast to LFTs who allocate enough that the corresponding distribution follows a power law (Fig. 18A, C). Overall, statistics illustrate that HFTs vary less in terms of trading strategies than LFTs. Classifying traders into these two distinct groups is therefore justified.

Modelling based on microscopic evidence. Focusing again on HFTs, we here construct a mathematical model that reflects the microscopic empirical evidence described heretofore. Let $N \gg 1$ denote the number of HFTs. The model's assumptions are:

1. *Order and volume.* Every HFT submits a single order at a time with a single unit of volume (Fig. 18A, B).
2. *Liquidity provision.* Every HFT keeps both bid and ask orders to play the role of a liquidity provider (Fig. 16).
3. *Frequent price updates.* HFTs frequently update their prices by successive order submissions and cancellations. This implies that the price trajectory is approximately continuous except at the times of transactions. The continuous Markov stochastic processes for price trajectories are modelled as an Itô process (i.e., a Gaussian stochastic process) [197].
4. *Trend following.* HFTs exhibit the trend-following behaviour in accordance to the empirical laws in Eq. (25) and Eq. (26) (see also Fig. 17). For simplicity, we assume the uniformity of model parameters in these equations such that $c_i = c$, $\Delta p_i^* = \Delta p^*$, and $\sigma_i^2 = \sigma^2$ for all i .
5. *Spread.* The buy-sell spread of the i th HFT is defined by $\hat{L}_i = \hat{a}_i - \hat{b}_i$. Because the probability density functions of buy-sell spreads have a single peak (insets in Fig. 16A–C), the spread is assumed unique to the HFT and constant, that is, $\hat{L}_i = L_i$. This assumption implies that the mid-price $\hat{z}_i = \frac{1}{2}(\hat{b}_i + \hat{a}_i)$ is sufficient to characterise the i th HFT.

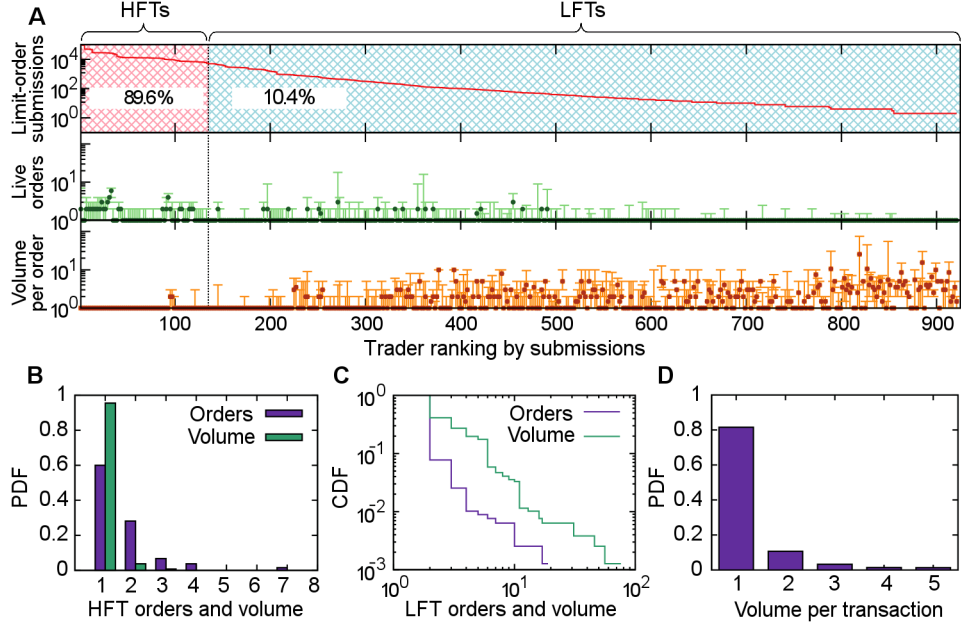


Figure 18: How HFTs and LFTs differ. **A**, Plots are an overview of limit-order submissions (top), live orders (middle), and volume assigned per order (bottom), according to trader ranking by submissions. Traders submitting more than 2,500 limit orders weekly are defined as HFTs; there have been 135 such traders in the available dataset, responsible for almost 90% of total submissions. **B**, Probability density functions reveal that HFTs maintain only a few live orders at a time, and that the volume per order is overwhelmingly one unit. **C**, Cumulative distribution functions reveal that LFTs use a wider range of volumes per order, suggesting a more diverse set of strategies than the one used by HFTs. **D**, Probability density function of the volume per transaction for orders that get filled shows that 81.5% of transactions are one-to-one, while the volume is less than five units for 98.2% of transactions.

Source: Reprinted figure from Ref. [192] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

The empirical probability density function of the buy-sell spread $\rho(L)$ is measurable from the available dataset using the relationship $\rho(L) = \frac{1}{N} \sum_{i=1}^N \delta(L - L_i)$, and thus describes the order-book distribution.

Based on the listed assumptions, we model the microscopic HFT dynamics in the absence of transactions as the trend-following random walks (Fig. 19A)

$$\frac{d\hat{z}_i}{dt} = c \tanh \frac{\Delta \hat{p}}{\Delta p_i^*} + \sigma \hat{\eta}_i^R, \quad (27a)$$

where the white noise $\hat{\eta}_i^R$ is independent of the white noise $\hat{\eta}_j^R$ for $j \neq i$. This is the minimal Itô process [197] satisfying the empirical relations we have described so far.

At the instant of price matching (Fig. 19B)

$$\hat{a}_j = \hat{b}_i, \quad i \neq j, \quad (27b)$$

the pair of HFTs i and j resubmit their prices far from the transaction price (Fig. 19C)

$$\hat{a}_j^{\text{pst}} = \hat{a}_j + \frac{L_j}{2}, \quad \hat{b}_i^{\text{pst}} = \hat{b}_i - \frac{L_i}{2}, \quad (27c)$$

where \hat{a}_j^{pst} and \hat{b}_i^{pst} are the updated prices after the transaction. These transaction conditions can be rewritten as

$$|\hat{z}_i - \hat{z}_j| = \frac{L_i + L_j}{2} \implies \hat{z}_i^{\text{pst}} = \hat{z}_i - \frac{L_i}{2} \text{sgn}(\hat{z}_i - \hat{z}_j), \quad (27d)$$

where the sign function $\text{sgn}(x)$ is defined by $\text{sgn}(0) = 0$, $\text{sgn}(x) = 1$ for $x > 0$, and $\text{sgn}(x) = -1$ for $x < 0$. The market price \hat{p} and the trend signal $\Delta\hat{p}$ at a transaction instant are updated with the post-transaction values

$$\hat{p}^{\text{pst}} = \hat{z}_i - \frac{L_i}{2} \text{sgn}(\hat{z}_i - \hat{z}_j), \quad \Delta\hat{p}^{\text{pst}} = \hat{p}^{\text{pst}} - \hat{p}. \quad (27e)$$

We note that the transaction condition in Eq. (27d) and the resubmission rule in Eq. (27e), respectively, bear mathematical resemblance to the contact condition and the momentum-exchange rule in conventional kinetic theory. This analogy will be revisited again later to formulate a statistical-physics description of financial markets. We also note that one-to-one transactions are the basic interaction mode between HFTs (Fig. 18D), which is consistent with binary collisions.

In statistical physics, an appropriate separation of spatio-temporal scales is often used to formulate successful micro-macro theories; an example is the enslaving principle in Haken's synergistics [198]. Here we introduce the centre of mass as a key macroscopic variable

$$\hat{z}_{\text{CM}} = \frac{1}{N} \sum_{i=1}^N \hat{z}_i, \quad (27f)$$

which is expected to play the role of a slow variable in the 'thermodynamic limit' when $N \rightarrow \infty$. Indeed, the diffusion of the centre of mass turns out

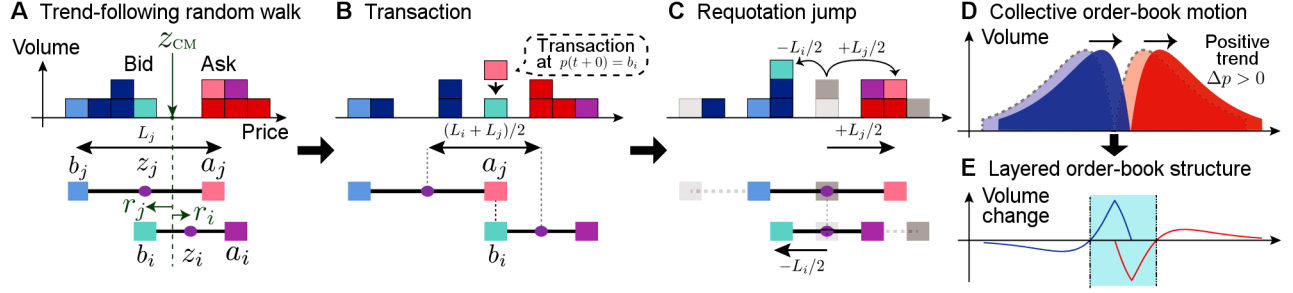


Figure 19: Schematic of the microscopic HFT dynamics and the corresponding order-book dynamics. **A**, Trend-following random walks take place in the absence of transactions. **B**, Transactions occur at the instance of price matching $\hat{a}_j = \hat{b}_i$. **C**, After a transaction between a pair of HFTs takes place, they resubmit their prices far from the transaction price. **D**, Trend-following random walks induce the collective order-book motion. **E**, Collective order-book motion consistently causes the layered order-book structure. Source: Reprinted figure from Ref. [192] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

to be slow (i.e., proportional to N^{-1} in the absence of trend following [192]), confirming the appropriateness of this particular variable selection.

A complete set of the system variables is given by $\hat{\Gamma} = (\hat{z}_1, \dots, \hat{z}_N; \hat{z}_{CM}, \hat{p}, \Delta\hat{p})$, with $\hat{Z} = (\hat{z}_{CM}, \hat{p}, \Delta\hat{p})$ being the subset of macroscopic variables. $\hat{\Gamma}$ can be regarded as a phase point in the space \mathcal{S} such that $\hat{\Gamma} \in \mathcal{S} = \prod_{i=1}^{N+3} (-\infty, \infty)$. We have thus defined a Markovian stochastic processes whose dynamics is characterised by the set of Eqs. (27).

The order-book dynamics associated with the described microscopic model (Fig. 19A–C) relates the trend-following behaviour to the layered order-book structure found in Ref. [185]. Specifically, trend following induces the collective motion of orders (Fig. 19D), which in turn accounts for the order book’s layered structure (Fig. 19E).

4.6. Solving the microscopic model via kinetic theory

Having presented a model of the decision-making process among HFTs based on empirical microscopic evidence, we proceed to solve this microscopic model in order to understand its macroscopic behaviour. Mathematically, we are dealing with a high-dimensional stochastic dynamical system exhibiting a structure similar to the Hamiltonian dynamics. We therefore rely on the methods of statistical physics in general, and kinetic theory in particular, to crack the problem.

Financial Liouville equation. Here, we start by briefly reviewing the conventional kinetic theory whose goal is to reduce the original high-dimensional dynamical system composed of N particles to a few-dimensional dynamical system. The microscopic dynamics (Fig. 20A) is generally characterised by Newton's equation of motion in $6N$ dimensions, which is mathematically equivalent to the Liouville equation in analytical mechanics [199, 200]. The Liouville equation reduces to the Boltzmann equation via the Bogoliubov-Born-Green-Kirkwood-Yvon (BBGKY) hierarchy [201] by applying the mean-field approximation called molecular chaos, thus offering a mesoscopic description of the system (Fig. 20B). This reduction is powerful in the sense that the original $6N$ -dimensional dynamics is approximated by a 6-dimensional dynamics. Further reduction of the dynamics yields the Brownian motion of a tracer particle [202, 203] as the macroscopic description of the system (Fig. 20C).

We have retraced the steps of the conventional kinetic theory with financial markets and the HFT behaviour in mind. From Eqs. (27), we derive the financial Liouville equation (equivalently, the Chapman-Kolmogorov equation [197] or the master equation [202]) as the time-evolution equation for the N -body probability density function $P_t(\Gamma)$

$$\frac{\partial P_t(\Gamma)}{\partial t} = \mathcal{L}P_t(\Gamma) \quad (28)$$

with an appropriate linear operator \mathcal{L} called the Liouville operator. We note that this equation is derived as an identity without any approximation and is equivalent to the original stochastic model described by Eqs. (27). This is an exact starting point for our statistical-physics theory.

Financial BBGKY hierarchy. While the financial Liouville equation (28) is exact, it cannot be solved analytically because it describes genuine many-body dynamics of a complex system. Here, we reduce this dynamical equation following the ideas behind the BBGKY hierarchy, which was historically invented for the purpose of a systematic derivation of the Boltzmann equation from microscopic particle dynamics. The derivation is based on reducing the dimension of the original high-dimensional dynamics by integrating out all variables except for a few dominating ones.

Because the derivation following the BBGKY hierarchy is long and technical [192], we only present the final results. Let us first define prices relative to the centre of mass, $\hat{r}_i = \hat{z}_i - \hat{z}_{\text{CM}}$, and the corresponding one-body and

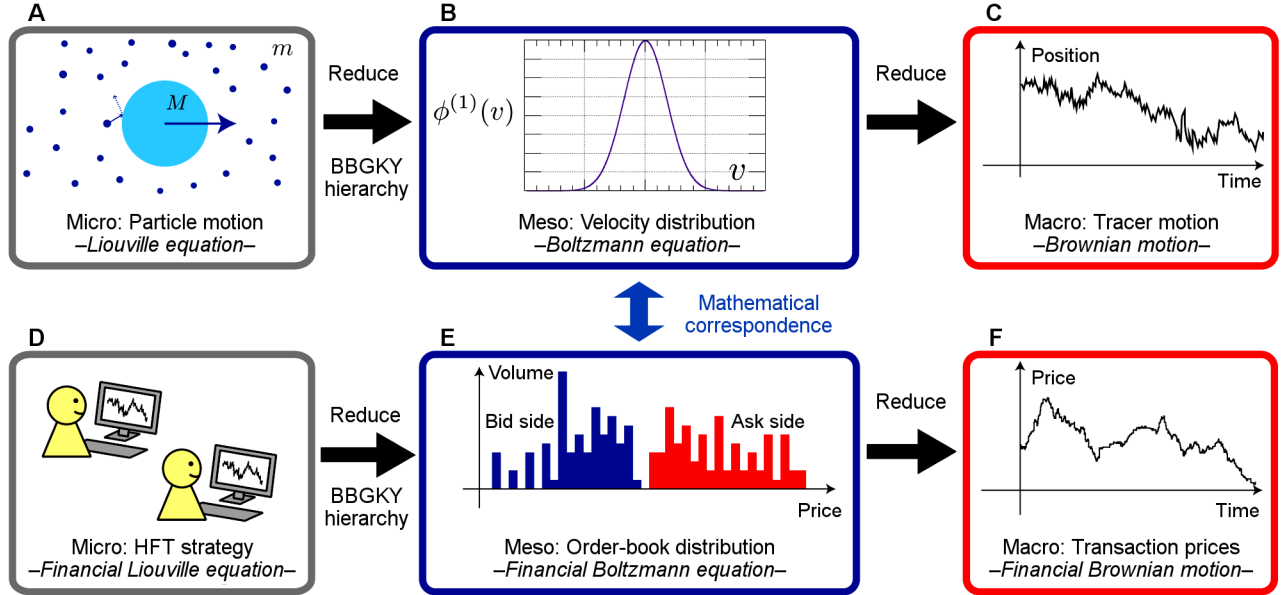


Figure 20: Schematic representation of the correspondence between kinetic theories for physical and financial Brownian motion. **A**, The starting point of physical kinetic theory is the Liouville equation for microscopic particle dynamics. **B**, The Liouville equation is reduced via the BBGKY hierarchy to the Boltzmann equation, which can be interpreted as the mesoscopic description of the dynamical system. **C**, The Brownian motion, described by the Langevin equation, is derived after further coarse-graining to yield the system’s macroscopic dynamics. **D**, The microscopic dynamics of financial markets is driven by the decision-making process of individual HFTs. This is captured by the financial Liouville equation (28), which is mathematically equivalent to the trend-following random-walks model in Eqs. (27). **E**, By reducing the microscopic market dynamics via the BBGKY hierarchy, we obtain the order-book dynamics in terms of the financial Boltzmann equation (31). The financial Boltzmann equation can be seen as the mesoscopic description of financial markets. **F**, Further coarse-graining reveals the market-price diffusion in the form of the financial Langevin equation (34).

Source: Reprinted figure from Ref. [192] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

two-body distributions $\phi_L(r)$ and $\phi_{LL}(r, r')$. Here, $\phi_L(r)$ is the probability density function of the relative mid-price r for an HFT with the spread L , while $\phi_{LL}(r, r')$ is the joint probability density function of a pair of HFTs with the spreads L and L' .

By integrating all but one variable out of Eq. (28), we obtain the lowest-

order BBGKY hierarchical equation

$$\frac{\partial \phi_L(r)}{\partial t} \approx \frac{\sigma^2}{2} \frac{\partial \phi_L(r)}{\partial r^2} + N \sum_{s=\pm 1} \int_0^\infty dL' \rho(L') [J_{LL'}^s(r + sL/2) - J_{LL'}^s(r)] \quad (29a)$$

$$J_{LL'}^s(r) = \frac{\sigma^2}{2} |\tilde{\partial}_{rr'}| \phi_{LL'}(r, r') \Big|_{r-r'=s(L+L')/2}, \quad (29b)$$

where $|\tilde{\partial}_{rr'}|f = |\partial f/\partial r| + |\partial f/\partial r'|$. The second term of the right-hand side of Eq. (29a) represents the effect of transactions ($s = +1$ for bids and $s = -1$ for asks) and corresponds to the collision integral in physical kinetic theory. The three-body ‘collision’ integral is dropped here because it becomes irrelevant for a large N .

Financial Boltzmann equation. The BBGKY hierarchical equation (29) is a formalism that needs a closure, that is, further approximation is necessary to find the solution. We apply the ‘molecular chaos’ assumption, which is a standard mean-field approximation in kinetic theory

$$\phi_{LL'}(r, r') \approx \phi_L(r)\phi_{L'}(r'), \quad (30)$$

to obtain the financial Boltzmann equation

$$\frac{\partial \phi_L(r)}{\partial t} \approx \frac{\sigma^2}{2} \frac{\partial \phi_L(r)}{\partial r^2} + N \sum_{s=\pm 1} \int_0^\infty dL' \rho(L') [\tilde{J}_{LL'}^s(r + sL/2) - J_{LL'}^s(r)] \quad (31a)$$

$$\tilde{J}_{LL'}^s(r) = \frac{\sigma^2}{2} |\tilde{\partial}_{rr'}| \phi_L(r)\phi_{L'}(r') \Big|_{r-r'=s(L+L')/2}, \quad (31b)$$

which is closed in terms of $\phi_L(r)$. This equation can be analytically solved under an appropriate boundary condition for large N . Indeed, the leading-order steady solution is given by the tent function

$$\lim_{N \rightarrow \infty} \phi_L(r) = \frac{4}{L^2} \max \left\{ \frac{L}{2} - |r|, 0 \right\}. \quad (32)$$

Notably, even the next-to-leading order solution is accessible, which is necessary for detailed mean-field analyses.

The average order-book profile follows from the financial Boltzmann equation via a formula

$$f^A(r) = \frac{1}{N} \left\langle \sum_{i=1}^N \delta(r - \hat{r}_i) \right\rangle \approx \int_0^\infty dL \rho(L) \phi_L(r - L/2), \quad (33)$$

where the summation is approximated by the integral over all spreads L . Remarkably, the book profile is analytically derived for any spread distribution $\rho(L)$, suggesting that the microscopic model of the HFT behaviour is an analytically well-tractable model.

Financial Langevin equation. By additional coarse-graining, we obtain the financial Langevin equation

$$\Delta \hat{p}[T + 1] \approx c \hat{\tau}[T] \tanh \frac{\Delta \hat{p}[T]}{\Delta p^*} + \hat{\zeta}[T] \quad (34)$$

for the macroscopic dynamics of the studied system. Here, $\hat{\tau}[T] = \hat{t}[T + 1] - \hat{t}[T]$ is the time interval between the ticks T and $T + 1$, whereas $\hat{\zeta}[T]$ is a random noise term. The mean-field approximation permits obtaining all the statistics for $\hat{\tau}$ and $\hat{\zeta}$ analytically. The time interval $\hat{\tau}$ exhibits the exponential distribution with the mean time interval $\tau^* \approx L_\rho^*/(2N\sigma^2)$ and $L_\rho^* = 1/\int_0^\infty L^{-2}\rho(L)dL$, implying the Poissonian statistics asymptotically.

The dynamical characteristics of the financial Langevin equation (34) depend on two dimensionless parameters, $\tilde{c} = cL_\rho^*/(\sigma^2\sqrt{2N})$ and $\Delta \tilde{p}^* = \Delta p^*/(c\tau^*)$. Focusing on a regime in which $\tilde{c} \gtrsim 1$ and $\Delta \tilde{p}^* \lesssim 1$, called the marginal-to-strong trend following [192], the statistics of the price change $\Delta \hat{p}$ asymptotically obeys the exponential law

$$P(\geq |\Delta p|; \kappa) \sim e^{-|\Delta p|/\kappa} \quad \text{for large } |\Delta p|, \quad (35)$$

where $P(\geq |\Delta p|; \kappa)$ is a complementary cumulative distribution function with the decay length κ .

Numerical confirmation. The validity of the described analytical predictions, such as the order-book profile in Eq. (33) and the exponential price-change distribution in Eq. (35), can be directly checked by means of numerical simulations [192]. The results are encouraging; especially the order-book profile formula in Eq. (33) agrees with the numerical results very precisely.

Such a precise agreement might be somewhat counter-intuitive because mean-field approximations are generally expected to be valid only for high-dimensional spaces while the price space is one-dimensional. This counter-intuitive result can be understood from the viewpoint of the ‘collision rule’: low-dimensional spaces are special for physical systems in the sense that geometry restricts movements of particles after collisions. The mean-field approximation then fails because the same pair of particles collides many times and the two-body correlation persists. In the case of our microscopic model, the ‘collision rule’ pushes the limit-order prices of a pair of HFTs far from the market price, making successive ‘collisions’ between the same HFT pair highly unlikely for a large N . The two-body correlation thus quickly decays, which to a large degree validates the assumption of the molecular chaos in our kinetic formulation. In fact, the kinetic formulation might work better for some social systems than for physical systems because the continuity of paths is often unnecessary in social dynamics but strictly required in physical dynamics.

4.7. Consistency between theory and data

We have analysed the model for trend-following random walks in Eqs. (27) within the kinetic framework. It is time to check the model’s consistency against mesoscopic and macroscopic data.

First, we have measured the buy-sell spreads of each individual HFT and estimated the corresponding daily distribution (Fig. 21A). The buy-sell spread distribution is well approximated by the γ distribution

$$\rho(L) \approx \frac{L^3}{6L^{*4}} e^{-L/L^*} \quad (36)$$

with $L^* \approx 15.5 \pm 0.2$ tpip. This buy-sell distribution together with the general order-book formula in Eq. (33) implies that the average order-book profile is given by

$$f_A(r) \approx \frac{4e^{-\frac{3r}{2L^*}}}{3L^*} \left[\left(2 + \frac{r}{L^*} \right) \sinh \frac{r}{2L^*} - \frac{r e^{-\frac{r}{2L^*}}}{2L^*} \right]. \quad (37)$$

Of note is that our model addresses the dynamics of the best bid and ask prices of individual HFTs. The average order-book profile after normalisation (Fig. 21B) shows an excellent agreement with the theoretical curve from Eq. (37) without any additional fitting of model parameters.

Turning to the macroscopic perspective in terms of the time series of price changes, we have seen that the financial Langevin equation (34) predicts the




Figure 21: Theory is consistent with data. **A**, Daily buy-sell spread distribution for individual HFTs is closely approximated by the γ -distribution in Eq. (36). **B**, Daily average order-book profile composed of the best HFT prices agrees with the theoretical curve in Eq. (37). **C**, Two-hourly segmented price-change cumulative distribution functions for three different time periods follow the exponential law in Eq. (35) with the decay length κ being time-dependent. **D**, After scaling, the two-hourly segmented price-change cumulative distribution functions collapse onto a single exponential master curve in Eq. (35). **E**, Weekly segmented price-change cumulative distribution function exhibits fat tails with a power-law exponent α . **F**, Decay length cumulative distribution function exhibits a fat tail with a power-law exponent m such that $Q(\geq \kappa) \sim \kappa^{-m}$.

Source: Reprinted figure from Ref. [192] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

exponential law in Eq. (35). This prediction turns out to be consistent with the available dataset (Fig. 21C, D). The price-change distribution at a time scale of one tick indeed follows said exponential law with the decay length κ depending on the chosen time period (Fig. 21C). Re-scaling price changes by κ eliminates the time-period dependence, causing all data to collapse onto a single master curve (Fig. 21D).

Although price changes obey the exponential law at short time scales, a power law emerges as time scales get longer, which is in line with previous findings [157, 204, 205, 206, 207]. The price-change distribution for the studied week has fat tails, fitted with a power-law whose exponent is

$\alpha = 3.6 \pm 0.13$ (Fig. 21E). The power law emerges here as a superposition of 2-hourly exponential distributions $P^{2h}(\geq |\Delta p|; \kappa)$ such that

$$P^w(\geq |\Delta p|) = \int_0^\infty d\kappa Q(\kappa) P^{2h}(\geq |\Delta p|; \kappa) \propto |\Delta p|^{-m}, \quad (38)$$

where $P^w(\geq |\Delta p|)$ is the weekly price-change cumulative distribution function, and $Q(\kappa) \propto \kappa^{-(m+1)}$ is the decay-length probability density function with the exponent being estimated at $m = 3.5 \pm 0.13$ (Fig. 21F). The fact that $\alpha \approx m$ additionally confirms the consistency of the results.

4.8. Future outlook: towards market ecology

We have examined in detail the trend-following behaviour of HFTs at the time scale of one tick. This examination had yielded a microscopic model specified by Eqs. (27) and solved using the kinetic theory of statistical physics. Ref. [208] extends the analysis, offering direct evidence that the (exponential) moving-average technique—a common tool in the repertoire of ‘technical analysts’ or ‘chartists’—is applied in practice by HFTs. To achieve this feat, a regression relation inspired by Eq. (27a) is set up as follows

$$\Delta \hat{z}_i[T] = c_i \tanh \left(\Delta \hat{p}_{\text{Trend}}^{(j_i)}[T] + \alpha_i \right), \quad (39a)$$

where the trend signal includes time delays

$$\Delta \hat{p}_{\text{Trend}}^{(j_i)}[T] = \sum_{k=1}^{K_i} w_i[k] \Delta \hat{p}^{(j_i)}[T - k] + \sigma_i \hat{\epsilon}_i[T]. \quad (39b)$$

Here, c_i , α_i , $w_i[k]$, σ_i , and K_i are regression coefficients, $\hat{\epsilon}_i[T]$ is the white noise, and $\{\Delta \hat{p}^{(j_i)}[T - k]\}_k$ is a coarse-grained price time series defined as

$$\Delta \hat{p}^{(j_i)}[T - k] = \hat{p}[T - j_i(k - 1)] - \hat{p}[T - j_i k], \quad (39c)$$

where k , K_i , and j_i are respectively called the time lag, the maximum time lag of the i th HFT, and the coarse-graining parameter. The regression coefficients are estimated using iterative non-linear multiple-regression methods.

The weights $\{w_i[k]\}_{k=1, \dots, K_i}$ determine the nature of the moving average applied by HFTs. The results show that an exponential scaling law is satisfied

$$w_i[k] \approx d_i \exp \left(-\frac{k}{\tau_i} \right), \quad (40)$$

where parameters d_i and τ_i characterise the i th HFT. Among the examined HFTs, 85 % were shown to conform to this exponential moving average. Further subdivision of HFTs was possible according to their typical trend-following time scale; (i) short time-scale HFTs follow the trend for about 4 ticks (30 s), (ii) intermediate time-scale HFTs follow the trend for about 20 ticks (3 min), and (iii) long time-scale HFTs follow the trend for about 40 ticks (6 min). The remaining 15 % of HFTs use other trading strategies.

Identifying market strategies as described here is a first step in the direction of understanding market ecology [209, 210], that is, interactions among various trading strategies (e.g., how strategies contribute to market liquidity or price formation). If such interactions are precisely understood, designing market simulators for regulatory purposes becomes entirely plausible. Regulators could then plan interventions to enhance market liquidity and stability. The scarcity of microscopic data has so far precluded in-depth analyses of market ecology (see Refs. [211, 212, 213, 214, 215, 216, 217, 218] for notable attempts so far), but we firmly believe that the situation is changing and that market ecology is within our grasp.

5. Cooperation

Conspecific organisms mutually interact to a lesser or greater degree. Some species are highly individualistic, associating only for courtship and mating. Others are extremely social with overlapping adult generations, reproductive division of labour, cooperative care of young, and sometimes even a biological caste system. Sociality, or living in groups, implies a coexistence of two opposing forces: conflict over local resources and cooperation with neighbours [219]. These two opposing forces form the basis of evolutionary game theory whose aim is to understand the evolution and pervasiveness of cooperation in biological and social systems. More specifically, the goal is to answer how natural selection can favour costly, cooperative behaviours that benefit others.

Ever since the mathematical framework of game theory was applied to evolution [220], the research on cooperation attracted the attention of fields as varied as biology, psychology, economics, physics, and others [221, 222, 223, 34]. Such a variety brought together a plethora of unique perspectives on the problem of cooperation, giving rise to a reasonably good understanding of the origins and stability of cooperativeness. It is safe to say that we are at

a point at which the evolution of cooperation is much less of a puzzle than it used to be [224].

Hereafter, we take a look at some of the central tenets of evolutionary game theory, and then review the main interests and contributions of physics to this field. A particular focus is on networks that define the topology of interactions among humans, as well as on a gap between theoretical models and empirical facts. We conclude with ideas for reconciling the gap between theory and experiments while heeding the initiatives from relevant behavioural disciplines (e.g., psychology and behavioural economics) for better research practices.

5.1. Social dilemmas

Social dilemmas are situations in which the process of selection favours defection over cooperation while reducing population welfare compared to when everybody cooperates [225]. Many real-world social dilemmas fit the format of dyadic games in which a pair of players simultaneously choose between cooperation C or defection D . Depending on the choices made, the outcome is defined with the following four payoff-matrix elements

$$\begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} R & S \\ T & P \end{pmatrix}. \end{array} \quad (41)$$

This payoff matrix signifies that mutual cooperation generates reward R , whereas mutual defection generates punishment P for both players. Additionally, if one player defects and other cooperates, the former receives temptation T and the latter the sucker's payoff S . Payoff ordering determines the nature of the dilemma. For example, $T > R > P > S$ indicates the prisoner's dilemma, that is, the archetypal dilemma for studying the emergence of cooperation between selfish individuals [226]. Two other common dilemmas are stag hunt and snowdrift (also known as hawk-dove or chicken), obtained by setting $R > T$ and $S > P$, respectively.

The above-mentioned social dilemmas were considered static prior to the work of John Maynard Smith, who introduced the notion of repetitions (i.e., iterations) and thus laid the foundations of evolutionary game theory [220, 227, 228]. Traditionally set in populations in which all players have equal probability to interact with one another (i.e., well-mixed populations), the frequency x_i ($0 \leq x_i \leq 1$) of players resorting to strategy i is traced with the

differential equation [229]

$$\frac{dx_i}{dt} = x_i [\varphi_i(\mathbf{x}) - \Phi(\mathbf{x})] \quad (42)$$

known as the replicator equation, where \mathbf{x} is the vector of frequencies for all strategies satisfying $\sum_i x_i = 1$, $\varphi_i(\cdot)$ is the per-capita payoff attained by resorting to the i th strategy, and $\Phi(\cdot)$ is the average per-capita payoff. Because dilemmas differ in payoff ordering, they also reach different stationary points under Eq. (42). If a stationary point is stable, it is called an evolutionarily stable state or a Nash equilibrium. If a Nash equilibrium is monomorphic, that is, $x_i = 1$ and $x_j = 0$ for all $i \neq j$, then the i th strategy is called an evolutionarily stable strategy (ESS). In prisoner's dilemma, defection is ESS. In a snowdrift dilemma, however, both C and D strategies coexist in an evolutionarily stable state [230] because cooperating with a defector is still better than mutually defecting.

Social dilemmas that fit the format of dyadic games can be re-scaled in terms of the dilemma-strength parameters, one of which (D'_g) measures how lucrative defection is in the presence of a cooperator, whereas the other (D'_r) measures how hazardous cooperation is in the presence of a defector [231, 232, 233]. Precisely, the two parameters are defined by

$$D'_g = \frac{T - R}{R - P}, \quad (43a)$$

$$D'_r = \frac{P - S}{R - P}. \quad (43b)$$

In Eq. (43a), the positive value of D'_g increases as T increases relative to R , which facilitates defection by making the temptation payoff for defecting against a cooperator much larger than the reward payoff for mutual cooperation. In Eq. (43b), the positive value of D'_r increases as S decreases relative to P , which again facilitates defection, but this time by making the sucker's payoff for cooperating with a defector much more negative than the punishment payoff for mutual defection. The normalisation factor in both equations, $R - P$, works in the opposite direction (i.e., facilitates cooperation) through more generous reward for cooperators and more stringent punishment for defectors. When $D'_g, D'_r > 0$, then the payoff ordering of the prisoner's dilemma holds; if instead $D'_g < 0$ ($D'_r < 0$), then the payoff ordering of the stag-hunt (snowdrift) dilemma holds (Fig. 22).




Figure 22: Dilemma-strength parameters reflect the nature of social dilemmas that fit the format of dyadic games. Specifically, $D'_g, D'_r > 0$ indicate the prisoner's dilemma; $D'_g < 0, D'_r > 0$ indicate the stag-hunt dilemma; $D'_g > 0, D'_r < 0$ indicate the snowdrift dilemma; and $D'_g, D'_r < 0$ indicate no dilemma (in which case the dyadic game is called harmony).

The importance of the dilemma-strength parameters lies in the fact that they reduce the dimensionality of dyadic games from four (payoffs R , S , T , and P) to two. This is achieved by affine-transforming the payoff matrix in Eq. (41) into

$$\begin{matrix} & C & D \\ \begin{matrix} C \\ D \end{matrix} & \begin{pmatrix} 1 & -D'_r \\ 1 + D'_g & 0 \end{pmatrix} \end{matrix}. \quad (44)$$

Such a transform leaves the process of selection as specified by Eq. (42) unchanged [227]. Of note is that infinitely many four-payoff matrices can be mapped into a single two-parameter matrix. Consequently, all dyadic games that have the same dilemma strength, even if their payoff matrices wildly differ, are equivalent in terms of evolutionary outcomes.

The prisoner's dilemma is ubiquitous because by default it leads to defection. Cooperation can prevail only through various extensions of dyadic

games based on this dilemma. Such extensions are then said to be cooperation-promoting mechanisms. Kin and group selection, as well as direct, indirect, and network reciprocity are seen as general mechanisms that act as the promoters of cooperation [234]. Interestingly, numerical results show that even when dyadic games incorporate these cooperation-promoting mechanisms, evolutionary outcomes are predetermined by the dilemma-strength parameters [231]. Why would that be? One way to understand why the quantities D'_g and D'_r work for extended dyadic games is to recognise that such games can all be transformed and reinterpreted as standard (i.e., non-extended) dyadic games but with suitably adjusted payoff matrices [235]. In the case of direct reciprocity, for example, the payoff matrix in Eq. (41) is transformed into

$$\begin{array}{c} C \\ D \end{array} \begin{array}{cc} C & D \\ \left(\begin{array}{cc} \frac{1}{q}R & S + \frac{1-q}{q}P \\ T + \frac{1-q}{q}P & \frac{1}{q}P \end{array} \right) \end{array}, \quad (45)$$

where q is the probability of terminating play with a given individual (and $1 - q$ is the probability of continuing play with this individual). Using the same re-scaling as in Eq. (44), we get

$$\begin{array}{c} C \\ D \end{array} \begin{array}{cc} C & D \\ \left(\begin{array}{cc} \frac{1}{q} & -D'_r \\ 1 + D'_g & 0 \end{array} \right) \end{array}, \quad (46)$$

which shows that for a fixed q , the evolutionary outcome is determined by the dilemma-strength parameters. If the probability q is small enough that $D'_g < \frac{1-q}{q}$, then the original prisoner's dilemma turns due to direct reciprocity to a stag-hunt dilemma, and cooperation is an ESS.

A more parsimonious explanation for the success of the dilemma-strength parameters is to recognise that all five cooperation-promoting mechanisms (i.e., kin selection, group selection, direct reciprocity, indirect reciprocity, and network reciprocity) have one crucial feature in common; one way or another, these mechanisms enable positive assortment by which cooperative acts occur more often between cooperators than expected based on population averages [236, 237, 238]. Once the role of positive assortment is recognised, the affine transformation of the payoff matrix as specified by Eq. (44) must not interfere with such assortment, that is, the two must be compatible. This indeed is the case [239]. In other words, all five cooperation-promoting mechanisms are manifestations of positive assortment which itself is preserved by

the affine transformation that parametrises dyadic games in terms of the dilemma-strength parameters.

A limitation of the dilemma-strength parameters is that they can be defined solely for social dilemmas that fit the format of dyadic games. However, a generalisation to more complex social dilemmas is possible in terms of an efficiency deficit defined as the fitness difference between the socially optimal steady state that maximises individual fitness and the current evolutionarily stable state [233]. This definition implies that society evolving to a suboptimal equilibrium incurs an opportunity cost. If the opportunity cost is small (large), it can (cannot) be tolerated, and a societal change for the better is more difficult (easier) to accomplish.

5.2. Cooperation in networks with pairwise interactions

A population of players can be structured using graphs or networks such that vertices (i.e., nodes) represent players and edges (i.e., links) indicate pairwise interactions. In this picture, the usual well-mixed populations of evolutionary game theory are represented by complete networks in which all nodes are linked to one another. Structured populations, however, are specified to spatially constrain interactions, that is, prescribe who interacts with whom. In a square lattice, for example, players interact with their four or eight nearest physical neighbours (often called von Neumann and Moore neighbourhoods, respectively). Payoffs accumulated from such interactions are then used to update the lattice through either reproduction or imitation and learning, depending on whether biological or social evolution is of interest.

When the focus is on biological evolution, ‘death-birth’ updating is commonly applied, meaning that at each time step a random player is chosen to die, followed by the offspring of neighbours competing for the empty site in proportion to their fitness [240, 241, 242]. An alternative is ‘birth-death’ updating by which a player is selected for reproduction proportional to fitness, followed by offspring replacing a randomly chosen neighbour. For social-evolution scenarios, ‘imitation’ updating is used, meaning that at each time step a random player is chosen to decide whether to keep their current strategy or imitate one of the neighbours depending on the difference in fitness between the neighbour and the player.

Structured populations shot to fame through the work of Nowak and May [243] who observed that repeated games in a square lattice generate spatial chaos. An even more influential finding was that cooperators could

expand by forming clusters [244] that enable reaping the benefits of cooperation despite exploitation by defectors at cluster boundaries. This finding spurred further studies on whether cooperators survive, or even thrive, in different types of network structures. Beside lattices, random and scale-free networks were featured prominently [245, 246, 247, 248, 249]. The cited studies, inspired in large part by the result in Ref. [250], established that more heterogeneous networks provide the best conditions for the evolution of cooperation by securing that large cooperative clusters remain little exposed to defection at cluster boundaries. The promise of scale-free networks strongly boosting cooperation succeeded in attracting much attention in the field, but has subsequently been proven to lack robustness to theoretical model assumptions [251, 252].

The proliferation of research on evolutionary games in structured populations begat a search for general rules that explain the evolution of cooperation in various networks. Ref. [240] describes one such rule for weak selection. The term ‘weak selection’ refers to the idea that many different factors affect a player’s overall fitness, with the game under consideration being just another factor among the many. For this reason, a player is characterised by a baseline fitness that is large relative to payoffs earned throughout the game. Let the game be a variant of the prisoner’s dilemma in which $R = b - c$, $S = -c$, $T = b$, and $P = 0$ (with $b > c > 0$). It turns out that, under weak selection, cooperation is favoured in pairwise networks if the benefit of altruistic acts, b , divided by the cost, c , exceeds the average number of neighbours, k , or $b/c > k$. This simple rule closely resembles Hamilton’s rule, according to which kin selection favours cooperation if $b/c > 1/r$, where r represents the coefficient of genetic relatedness between individuals. The similarity of the two rules can be intuitively understood by considering that the average node degree is an inverse measure of social relatedness between players. A player’s fate is loosely bound to that of the neighbours if there are many neighbours, whereas the opposite is true if there are a few neighbours. Similarity notwithstanding, Ref. [253] argues that network reciprocity under the condition $b/c > k$ is fundamentally different from kin selection under the condition $b/c > 1/r$. Importantly, the network-reciprocity rule is rather robust to theoretical model assumptions [252].

Ref. [254] expanded the aforementioned research to obtain a general condition for the evolution of cooperation in any network under weak selection. Writing the condition using our notation, we have $\sigma R + S > T + \sigma P$, where $\sigma = \frac{-t_1+t_2+t_3}{t_1+t_2-t_3}$ is a structural coefficient that quantifies a network’s propensity

to support cooperation. The quantities t_1 , t_2 , and t_3 respectively denote the expected times at which the first, second, and third neighbours of an initial cooperator become cooperators. To revert back to the simple rule $b/c > k$, in addition to setting the payoffs to $R = b - c$, $S = -c$, $T = b$, and $P = 0$, it is necessary that the network is large enough in order for all nodes (whose average degree is k) to be sparsely connected.

The research on the evolution of cooperation in structured populations described so far has maintained a rigid assumption of node-based selection. This means that any given node is either a cooperator or a defector and acts as such towards its whole neighbourhood. In reality, it is crucial for many simple and complex organisms alike to differentiate between cooperative and defecting neighbours. By refocusing selection on links instead of nodes, a series of recent works has enabled examining situations in which the same node can cooperate with some neighbours and defect against others. The results show that with link-based selection, the frequency of cooperation can be high for a wide range of game setups [255]. A novel dynamic state has been observed between $b/c \approx \langle k \rangle$ and $b/c \gg \langle k \rangle$ in which cooperation and defection dynamically interchange with one another as a dominant strategy [256]. In mixed populations with both node- and link-based selection, cooperation either increases monotonically as the link-based selection becomes more prevalent [257], or there is a clear separation of roles by which node-based selection spawns cooperative clusters, while link-based selection protects these cluster from defectors [258].

When evolutionary games in structured populations are enriched with a third strategy on top of usual cooperation and defection, a commonly observed phenomenon is that of cyclic dominance [259]. The term ‘cyclic dominance’ refers to an intransitive relationship between objects A, B, and C by which A in some aspect dominates B, B dominates C, but C dominates A. The third strategy leading to cyclic dominance can be as simple as that of loners [260], who always stay out of the game, thus settling for a small payoff no matter whom they were supposed to interact with. Similar results are seen with exiters [261], who pay a small cost to find out if they are supposed to interact with cooperators or defectors. In the former case, exiters stay in the game and cooperate. In the latter case, they exit the game to receive a small payoff before getting exploited through defection. Yet another example of a strategy leading to cyclic dominance is that of hedgers [262], who also pay a small cost to find out if they are supposed to interact with cooperators or defectors, but then cooperate in the former case and defect

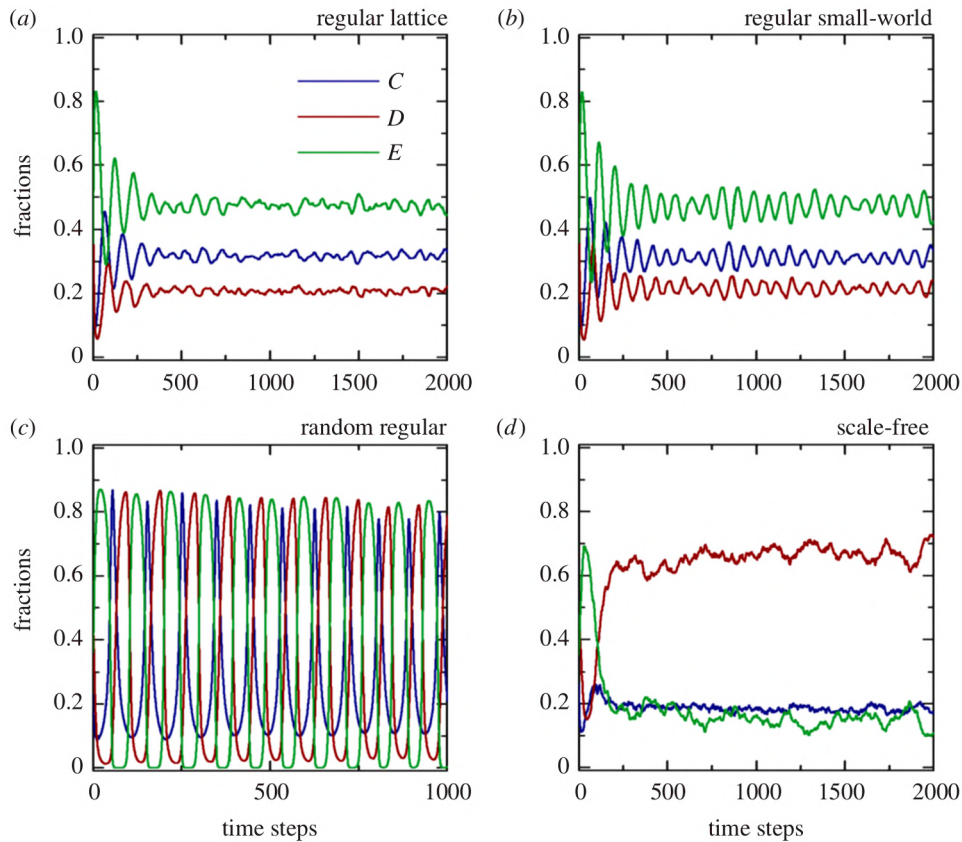


Figure 23: Evolutionary dynamics of cooperators, defectors, and exiters depending on network topology. Panel (a) shows that abundances of cooperators, defectors, and exiters oscillate locally in regular lattices. Panel (b) further shows that local oscillations are somewhat amplified by regular small-world networks. Panel (c) reveals that global oscillations occur in random regular networks. In scale-free networks, shown in panel (d), the presence of hub nodes turns oscillations into random fluctuations. Source: Reprinted figure from Ref. [261].

in the latter. Cyclic dominance is of interest from a dynamical point of view because evolutionary dynamics can greatly differ depending on the topology of the underlying network (Fig. 23).

To conclude, network structure may be a powerful cooperation-promoting mechanism in the prisoner's dilemma, but this is not the case in the snowdrift dilemma. Structure, surprisingly, decreases the frequency of cooperators relative to well-mixed populations [263, 230]. Instead of large, compact clusters

common in spatial prisoner’s dilemma games, clusters in spatial snowdrift dilemma games are small and filament-like [230] due to the fact that two interacting snowdrift players should adopt the opposite strategies to one another.

5.3. Cooperation in multilayer networks

Many complex systems can be seen as a network of networks. Organisms, for example, comprise gene regulatory networks on the sub-cellular scale, neuronal networks on the cellular scale, and vascular networks on the scale of cellular collectives (i.e., organs) [264]. Ecosystems comprise trophic networks and host-parasite networks in habitat patches accessible to individuals [265]. Infrastructure comprises many interdependent networks such as power, communications, transportation, etc. Interdependence in particular implies that processes occurring in one network may affect what happens in other networks to the point that small and seemingly irrelevant changes have unexpected and catastrophic consequences [266]. This possibility has sparked substantial interest in the robustness of networks of networks in general, as well as in many specific contexts [267, 268, 269, 270]. Interestingly, a wide variety social interactions fit a network of networks representation. For example, people mutually interact and transfer information both within and between online social networks. It is therefore natural to study the evolution of cooperation in interdependent networks [271].

A rigorous way of representing networks of networks is via the multilayer-network formalism [272, 273]. Keeping the discussion semi-formal, we can define a multilayer network as a quadruplet $M = (V_M, E_M, V, L)$, where V is a set of physical nodes, $L = L_1 \times \dots \times L_d$ is a set of layers comprising d elementary-layer sets L_1 to L_d , $V_M \subseteq V \times L$ is a set of state nodes encoding whether a physical node $v \in V$ is found in layer $l \in L$, and $E_M \subseteq V_M \times V_M$ is a set of intralayer and interlayer links. To exemplify, an elementary set could be an online social network (e.g., $L_1 = \{\text{Facebook, Twitter}\}$), while the other elementary set could be a region (e.g., $L_2 = \{\text{US, EU}\}$). Then the set of layers is $L = L_1 \times L_2 = \{(\text{FB, US}), (\text{FB, EU}), (\text{TWTR, US}), (\text{TWTR, EU})\}$. To indicate that the person v_1 accesses Facebook and Twitter from the US, we would write $(v_1, (\text{FB, US})) \in V_M$ and similarly $(v_1, (\text{TWTR, US}))$, which could be shortened to $(v_1, (1, 1))$ and $(v_1, (2, 1))$. The person v_1 is also an interlink between the two social networks $\{(v_1, (1, 1)), (v_1, (2, 1))\} \in E_M$. To indicate that the person v_2 is a Twitter user who travels between the US and the EU, we would write $(v_2, (2, 1))$ and $(v_2, (2, 2))$, automatically forming an

interlink $\{(v_2, (2, 1)), (v_2, (2, 2))\}$. Finally, if the person v_1 follows the person v_2 on Twitter, this can be represented as the intralink $\{(v_1, (2, 1)), (v_2, (2, 1))\}$ and the interlink $\{(v_1, (2, 1)), (v_2, (2, 2))\}$. Another intuitive example of a multilayer network could have a collection of cities as physical nodes; the first elementary layer could be transport-mode availability (e.g., high-speed railway station or airport), the second could be city size (e.g., less than 1 million inhabitants, between 1 and 5 million inhabitants, or more than 5 million inhabitants), and the third elementary layer could be a country. Instances of intralayer links would then be domestic railways or flights connecting the cities of the same size. Instances of interlayer links would be any international railways or flights, domestic railways or flights between the cities of different size, but also cities that possess both a high-speed railway station and an airport where transferring between the two transport modes is possible.

From the perspective of evolutionary game theory, the focus is on interdependence through the coupled player utilities or the flow of information between players, although other proposals have been made too [271]. Typically, a single player occupies only one of the available layers (but see Ref. [274]), while gameplay and strategy transfers take place only between players residing in the same layer. If the latter were not the case, and we dealt with truly interconnected networks (as opposed to just interdependent ones), then the social-dilemma game would essentially unravel in a single-layer network (albeit with two communities), meaning that the same cooperation-promoting mechanism would operate everywhere. Ref. [275] presents a game setup designed along these lines, that is, the prisoner's dilemma is played within and between two communities. In Ref. [276], the prisoner's dilemma is played within communities and the snowdrift dilemma between them.

The simplest and most common social-dilemma games in multilayer networks are those unfolding in two-layer networks (Fig. 24). Ref. [277] exemplifies such a game with coupled player utilities. Specifically, let x denote a player from layer 1 whose payoff is P_x , and similarly, let x' denote an interdependent player from layer 2 whose payoff is P'_x . Then the utility determining the course of evolutionary dynamics for both players is given by

$$U_x = U'_x = \alpha P_x + (1 - \alpha) P'_x, \quad (47)$$

where $0 \leq \alpha < \frac{1}{2}$. Furthermore, instead of playing the prisoner's dilemma with their first neighbours individually, the players x and x' participate in public goods games with all their first neighbours simultaneously. In a lattice,




Figure 24: Two-layer network for a social-dilemma game. Nodes in different layers represent different players, which is denoted by blue and red colours. Although players reside in different layers, there are interdependencies between them, which is denoted by the dashed interlinks. The topology of one layer, as indicated by solid intralinks, may differ from that of another layer, meaning that social connectivity is layer-specific. Source: Reprinted figure from Ref. [271].

the public goods game can be centred around the player x in layer 1 (x' in layer 2), but also around each of the first neighbours, meaning that the players x and x' participate in five public good games to collect their payoffs, P_x and P'_x , respectively. The gameplay rules are such that a cooperator contributes 1 unit to a pool, while a defector contributes nothing. The total contribution to the pool is multiplied by a return factor $r > 1$, and divided equally between players participating in the same public goods game. The results show that cooperation is strongly promoted in layer 1, but not in layer 2, due to a dampening effect that coupled utilities have on the exploitation of cooperators by defectors in the former layer.

Aiming to resolve the social dilemma in both layers, Ref. [278] redefined the coupled utilities

$$U_x = P_x + \alpha_x P'_x, \quad (48)$$

$$U'_x = P'_x + \alpha'_x P_x, \quad (49)$$

where $0 \leq \alpha_x, \alpha'_x \leq 1$ are directed interdependencies between the layers. These interdependencies are adaptive such that if the player x (x') earns a payoff P_x (P'_x) greater than some threshold E , that is, $P_x \geq E$ ($P'_x \geq E$), then the interdependency α_x (α'_x) is strengthened by an amount $\delta > 0$; if conversely the payoff falls short of the threshold, the interdependency is reduced by the same amount.

It turns out that under the described setup, a large threshold value E effectively keeps the two layers disconnected, meaning that there cannot be any synergy between them. Somewhat surprisingly, if the threshold value E is too small, suboptimal synergy is achieved. This is because small E allows defectors (alongside cooperators) to develop strong interdependencies, and then, if a defector in one layer gets to exploit some cooperators, then this defector can sustain their counterpart defector in the other layer. Only for the intermediate values of the threshold E is full synergy between the layers achieved (Fig. 25). In this case, predominantly cooperators create strong interdependencies which allows two interdependent cooperators to support one another. Such cooperators seed clusters of cooperation that later on take over the whole population in both layers (Fig. 25), thus successfully resolving the social dilemma at hand.

5.4. Cooperation in temporal networks

Complex networks have heretofore been used as if social interactions were static in time. Although this is a reasonable approximation in many circumstances, the ephemeral nature of human contacts eventually needs to be accounted for. After all, when two persons engage in an activity, this often happens in short bursts followed by periods of relative lull [279, 280, 281].

Temporal networks have emerged as a convenient tool for representing time-varying social interactions [282, 283, 284]. In such networks, any pair of momentarily disconnected nodes may get connected by the next time instant and vice versa. Interest in temporal networks stems from their ability to affect network-science fundamentals, among others, general dynamical processes [285, 286, 287], epidemiological dynamics [288, 289], and network controllability [290].

In the study of the evolution of cooperation in networked populations, network temporality has found a natural place in co-evolutionary models [291]. The term ‘co-evolution’ implies that beside the usual cooperative trait, one or more other traits evolve in parallel. This could be, for instance, a homophilic trait such that cooperative individuals tend to connect with other cooperative individuals or, at least, shun connections with defecting individuals. Be it homophily or some other psycho-social mechanism, a consequence is that network topology changes over time [292, 293, 294, 295, 296]. Psycho-social mechanisms, however, imply an active screening for suitable contacts, whereas network temporality in the real world is oftentimes more serendipitous. This begs the question of how temporal contact networks, exogenous




Figure 25: Snapshots of co-evolutionary dynamics in a two-layer lattice. Interlayer links occupied by cooperator-cooperator pairs acts as seeds for the growth of cooperative clusters. Panels (a) show cooperation frequency in the two lattice layers, whereas panels (b) show directional interdependency strength between the two lattice layers. From left to right, the panels display Monte Carlo Steps 0, 100, 300, 500, and 9999, respectively. Source: Reprinted figure from Ref. [278].

of psycho-social mechanisms, affect the evolution of cooperation.

Early work in the described context has indicated that temporal networks may be favouring selfish behaviour [297]. Using a temporal network of $N = 100$ contacts recorded every five minutes over a period of six months (resulting in 41,291 snapshot graphs), it was shown that when snapshots are aggregated over a short time period of $\Delta t = 1$ h, then much lower cooperation frequencies ensue compared to longer aggregation periods of $\Delta t > 1$ wk. Here, aggregation means taking all snapshot graphs over the period Δt , and if two nodes i and j interacted in any of the snapshots, then these nodes are assumed to be momentarily connected; otherwise, the nodes are disconnected (Fig. 26A). The time period Δt is a facet of network temporality, where

small (large) Δt values mark frequent (infrequent) changes in topology. As $\Delta t \rightarrow \infty$, the network becomes fully aggregated (i.e., static). Interestingly, randomising the time ordering of snapshots improves cooperativeness. Because such randomisation effectively removes the aforementioned burstiness of human interactions, it would seem that the bursts of activity, in particular, disfavour the evolution of cooperation.

In contrast to the outlined early work, the current state of the art [298] paints a more nuanced picture of cooperativeness in temporal networks. This is achieved by considering an additional facet of network temporality; how fast evolutionary dynamics is relative to network-structural dynamics. A parameter g quantifying the second facet of temporality is defined as the number of evolutionary-game rounds that take place during the time period Δt (Fig. 26B). The increasing values of the parameter g improve cooperativeness in temporal networks even beyond what is possible in static networks (Fig. 27). The improvement occurs despite the unfavourable effects of burstiness.

The numerical results in Ref. [298] point to a threshold for the outbreak of defection that reaches a maximum for intermediate values of the aggregation parameter Δt . This result can be understood via activity-driven modelling [299], which shows that defectors successfully spread if the following condition is satisfied

$$\frac{\lambda k}{\mu} \geq D^*, \quad (50)$$

where λ (μ) is the average probability of a cooperator (defector) turning into a defector (cooperator) in the next round. The quantity $k = 2l\langle a \rangle$ is the average degree given in terms of the average number of links l that an active temporal-network node randomly creates in the current time step, as well as in terms of the average activity $\langle a \rangle$. To clarify, in activity-driven models, each node is assigned a probability a_i of being active in a particular time step; active nodes create l random links to other (active or passive) nodes, while also being able to receive additional links from other active nodes. The quantity on the right-hand side of the invasibility condition in Eq. (50) is

$$D^* = \frac{1}{1 + \sqrt{1 + \frac{\text{Var}(a)}{\langle a \rangle^2}}}. \quad (51)$$

This relationship shows that the threshold D^* should decrease as the aggregation period Δt gets very short because then temporal contact networks are




Figure 26: Evolutionary games in temporal networks. **A**, Social interactions change from one time instant to another (left). This is represented using temporal networks after aggregating all social interactions over the time period of length Δt (right). The longer the Δt is, the coarser the picture of social interactions. When Δt is large, we get only the fully aggregated (i.e., static) interaction network. **B**, The aggregation-length parameter Δt captures one facet of network temporality. Another facet is captured by a parameter g , quantifying how fast evolutionary dynamics compared to network-structural dynamics. The parameter g is defined as the number of evolutionary-game rounds between any two consecutive changes in the network topology. Source: Reprinted figure from Ref. [298] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

sparse (see Fig. 26A for $\Delta t = 3$), which makes the average activity $\langle a \rangle$ also small. A sparse contact network removes the benefits of network reciprocity. Nodes end up playing pairwise repeated prisoner's dilemma games in which a likely Nash equilibrium is defection (depending on game payoffs and repetitions). Furthermore, the threshold D^* should decrease as the aggregation period Δt gets very long because then temporal contact networks are highly heterogeneous (Fig. 26A for $\Delta t = 12$), which makes the variance $\text{Var}(a)$ large.

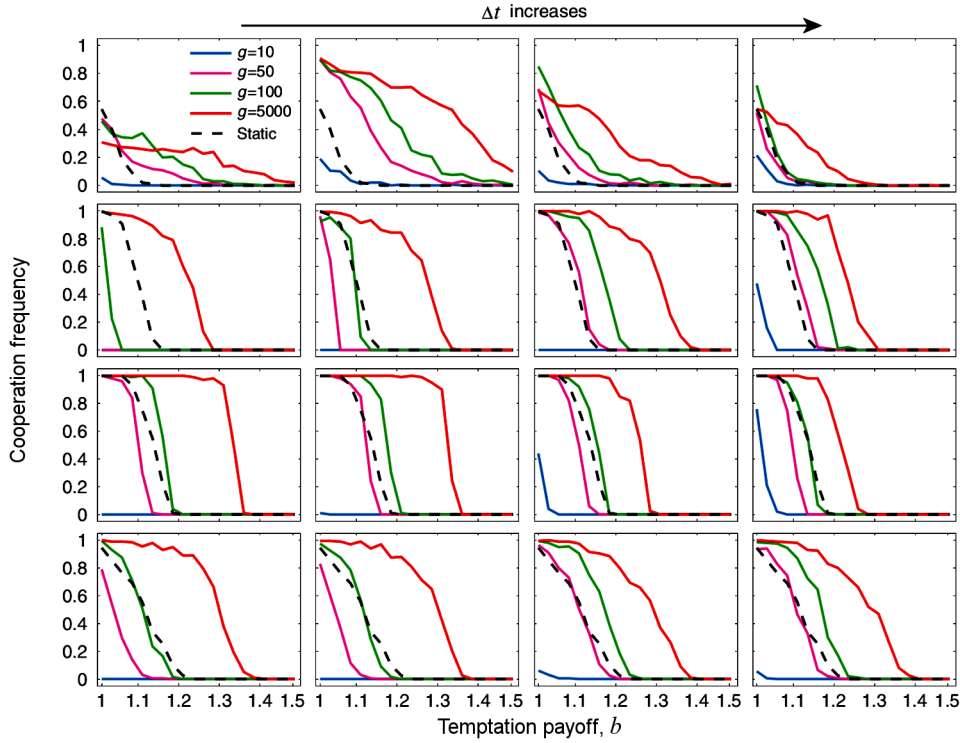


Figure 27: Temporal networks can promote the evolution of cooperation beyond their static counterparts. Panels show the cooperation frequency as a function of the temptation payoff for several values of the parameter g . Each row of panels corresponds to one real-world temporal network: 1st row is a contact network of attendees at a scientific conference, 2nd and 3rd rows are contact networks of students at a high school in Marseilles, France, in 2012 and 2013, and 4th row is a contact network of workers in an office building in France. Irrespective of the network or the value of the aggregation parameter Δt , there is always some g value for which the cooperation frequency in the temporal network is larger than in the corresponding fully-aggregated (i.e., static) network.

Source: Reprinted figure from Ref. [298] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

A heterogeneous contact network undergoes big topological transformations from one time step to another. Any such transformation can destabilise cooperative clusters, thus making cooperators more vulnerable to defectors.

In summary, temporal networks may promote cooperation beyond what is possible in the corresponding fully aggregated (i.e., static) networks, but only if the evolutionary-game dynamics is fast relative to the network-structural dynamics, that is, the parameter g is large. An additional important result is

that temporal networks are most resistant to defection for the intermediate values of the aggregation period Δt . These positive effects may, however, be nullified by the burstiness of human interactions in instances when the two facets of temporality, g and Δt , are unfavourable.

5.5. Cooperation in networks with higher-order interactions

As illustrated thus far, studying the evolution of cooperation in graphs or networks has a long tradition. Intriguingly though, cooperation in groups that are themselves embedded in a higher-order network structure has remained an open question until recently [300]. This is not to say that the question has been entirely ignored; special cases have been examined and have, in fact, left their mark on the field [301, 302]. These early works would select a focal node and follow this node’s pairwise links to determine which other nodes form a group. Thus determined group members would then participate in the same public goods game. A problem here is that determining group members based on pairwise links is rather unsatisfactory. Pairwise links are, after all, supposed to denote interactions between node pairs, not node groups, leading to many natural questions. Can every node be focal or, if not, how do we select focal nodes? Should a group extend only to the focal node’s first neighbours or should it include second, third, or even more distant neighbours? Does the higher-order network of groups that emerges by clumping together focal nodes and their neighbours have the most general structure possible or are there substantial limitations to what is achievable?

Among the most straightforward generalisations of ‘classical’ networks to higher-order ones is to allow more than two nodes to be connected via the same link, in which case the term hyperlink becomes customary [303]. The resulting higher-order network is often referred to as a hypergraph. Formally, a hypergraph H is a pair of sets $H = (N, L)$, where $N = \{n_1, \dots, n_n\}$ is a set of nodes and $L = \{l_1, \dots, l_l | l_j \subseteq N\}$ is a set of hyperlinks. Because hyperlinks themselves are subsets of the set N , a hyperlink’s number of elements (i.e., its cardinality) is a well-defined concept. The cardinality, $|l_j| = g$, is usually called the order of the hyperlink and is used in generalising the idea of the node degree. Specifically, if k_i^g denotes the number of order- g hyperlinks containing the i th node, then this node’s hyperdegree is $k_i = \sum_{g=g_{\min}}^{g_{\max}} k_i^g$, whereas the average hyperdegree is $\langle k \rangle = \frac{1}{n} \sum_{n_i \in N} k_i$. These definitions permit introducing two types of heterogeneity into hypergraphs. First, hyperdegrees can differ from one node to another, while all links have the same order g . Second, and more generally, links of multiple orders can

intermix in such a way that k^g/k is the probability that a randomly chosen hyperlink is of the order g , where $g_{\min} \leq g \leq g_{\max}$. This means that a node of hyperdegree k on average belongs to k^g order- g hyperlinks. For a hypergraph to be uniform, the node hyperdegree must be the same across all nodes and the hyperlink order must be the same across all hyperlinks.

Hypergraphs are an ideal setting to study cooperativeness in groups of individuals in a general situation when each individual potentially belongs to more than one group. An individual is represented by a hypergraph node, whereas a group comprises all nodes connected via a single hyperlink. Because every group has two or more individuals, it is natural to consider cooperation in the public goods game, that is, a social-dilemma game that is the multiplayer generalisation of the prisoner's dilemma [304, 233]. The usual game-rules apply. Cooperators pay the cost c that is pooled, multiplied by a return factor, and then split equally among all game participants, even if they are defectors and refuse to contribute to the pool. If we set $c = 1$ and denote with r the return factor divided by the number of game participants, then the per-capita payoff of a cooperator is $\pi_C = \nu_C r - 1$, whereas that of a defector is $\pi_D = \nu_C r$, where ν_C denotes the number of cooperators. It is clear that $\pi_C < \pi_D$ for $\nu_C > 0$, but if nobody cooperates, nobody gets any return either. Therein lies the dilemma. Strategy selection proceeds such that a focal node n_i is chosen randomly, as is a hyperlink l_j to which this node belongs. Then all members of the hyperlink l_j (i.e., $n_{i'} \in l_j$) play one public goods game in each of the hyperlinks they belong to. This provides the average payoff that nodes $n_{i'}$ earn per game played. The focal node finally adopts the strategy of the best performing neighbour with the probability $\frac{1}{\Delta} (\max \pi_{i'} - \pi_i)$, where Δ is a normalising quantity equal to the absolute maximal payoff difference over all the possible strategies.

It is a widely known result that in a well-mixed population, for cooperation to evolve in a public goods game, the condition $r > 1$ must be satisfied (e.g., see Refs. [233, 300]). For $r < 1$, defection prevails. Do hypergraphs help to promote cooperation in the sense of relaxing the condition $r > 1$? Although the answer to this question is technically positive, the cooperation-promoting effect of hypergraphs is limited. Starting with uniform hypergraphs, for a given order g , there is a critical number of hyperlinks l_c that is needed to guarantee the existence of a giant connected component. When the actual number of hyperlinks l is of the order of l_c , non-zero cooperation does appear for the critical return factor $r_c < 1$, but full cooperation is possible only for $r > 1$ (Fig. 28A). In fact, the critical return factor tends rather

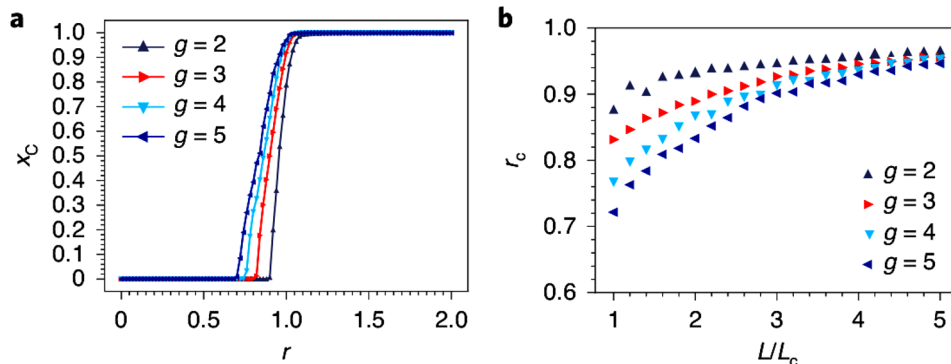


Figure 28: Cooperation in uniform hypergraphs. Panel **a** shows that unlike in well-mixed populations, non-zero cooperation is possible for return-factor values $r_c \leq r < 1$. However, full cooperation evolves only for $r > 1$. In these simulations, the number of hyperlinks equals the critical number that guarantees hypergraph connectedness, $l = l_c$. Panel **b** reveals that as the number of hyperlinks l increases above the critical value l_c , the critical return factor r_c necessary for non-zero cooperativeness tends to unity. A quick rate of convergence indicates that even relatively sparse hypergraphs mimic well-mixed populations. Source: Reprinted figure from Ref. [300].

quickly to unity, $r_c \rightarrow 1$, as the actual-to-critical hyperlink ratio, l/l_c , increases (Fig. 28B), indicating that even relatively sparse hypergraphs mimic the well-mixed population when it comes to the evolutionary dynamics of cooperation. It is important to note that, for a constant l/l_c , the increasing value of g makes hypergraphs sparser, which fully explains the dependence of the results in Fig. 28 on the hyperlink order. It is furthermore of interest that introducing moderate hyperdegree heterogeneity has no impact on the results [300]. Only in scale-free hypergraphs new patterns emerge, but unlike in classical networks, establishing cooperation becomes more difficult because large-hyperdegree nodes connect hypergraph parts that would otherwise be disconnected, thus effectively counteracting sparsity [300].

The other type of hypergraph heterogeneity, that is, order heterogeneity, is not so much of interest in the context of promoting cooperation as much as there are insights to be gained about the performance of collaborative groups [305]. Ref. [300] assumes that the return factor depends on hyperlink order, $r = r(g)$. If this is the case, the usual difference between the expected

per-capita payoff of cooperators and defectors, $\bar{\pi}_C - \bar{\pi}_D = r - 1$, changes to

$$\bar{\pi}_C - \bar{\pi}_D = \sum_{g=g_{\min}}^{g_{\max}} \frac{k^g}{k} [r(g) - 1], \quad (52)$$

where for each hyperlink order g , we take into account the corresponding probability (given by the fractional multiplier) and the corresponding outcome (given by the square-bracket multiplicand). It is natural to assume that the return factor consists of two parts, a benefit part $\alpha g^{(\beta-1)}$ that increases with the hyperlink order due to synergies of working with collaborators, and a cost part $\exp[-\gamma(g-1)]$ that decreases with the hyperlink order due to difficulties of coordinating large groups. We thus have

$$r(g) = \alpha g^{(\beta-1)} e^{-\gamma(g-1)}, \quad (53)$$

where the parameters α , β , and γ need to be estimated. Such estimation is doable if the return factor is extracted from data using the following two-step procedure:

1. Set $r(g) \propto \frac{k^g}{k}$ based on the intuition that a node hyperdegree distribution should align with potential returns, that is, the most probable hyperlink orders are also the ones that return the most.
2. Set the expected per-capita payoff for cooperating and defecting to be the same, that is, $\bar{\pi}_C = \bar{\pi}_D$. This guarantees that cooperators and defector coexist, as is often the case in real-world collaborations.

The two-step procedure is illustrated in Ref. [300] on 577,886 papers published by the American Physical Society (APS) from 1904 to 2015. Specifically, a hypergraph is constructed for each of 13 society journals such that nodes represent scientists publishing in a chosen APS journal, whereas hyperlinks represent articles. The hypergraph then provides the probabilities k^g/k , while the condition $\bar{\pi}_C = \bar{\pi}_D$ is used to fix the proportionality constant between $r(g)$ and k^g/k . Thus determined values of $r(g)$ are finally used to fit the parameters of Eq. (53). Several valuable lessons about collaboration in physics follow. First, collaborations between two to three scientists offer the best cost-benefit performance in most journals (Fig. 29A). Collaborating is therefore beneficial, but the costs of coordinating larger groups become substantial rather quickly. Exceptions are *Physical Review Series I*, which




Figure 29: Optimal size of collaborative groups in physics. Panel **a** shows the relationships between the return factor and the hyperlink order as implied by 13 hypergraphs constructed for the same number of the American Physical Society (APS) journals. Nodes are individual authors and hyperlinks are articles. As the hyperlink order (and thus collaborative group size) increases, at first the return factor increases as well, signifying benefits from distributing research work. Thereafter, the return factor decreases, signifying costs of coordinating large collaborations. Panel **b** shows that cost and benefit parameters follow approximately linear correlation, which ultimately keeps the optimal size of collaborative groups between two and five.

Source: Reprinted figure from Ref. [300].

includes publications up to the year 1913 when publishing alone was still very much feasible, as well as *Physical Review Applied* and *Physical Review X*, which show that applied and interdisciplinary research profit from assembling larger collaborative groups. The benefit parameter β and the cost parameter γ approximately exhibit a linear correlation (Fig. 29B). When benefits increase with group size, costs do too, keeping the optimal group size within two to five scientists across all modern APS journals.

5.6. Empirical facts about human cooperation

Studies on the evolution of cooperation, especially those based on modelling, have proliferated over the past two decades. To exemplify, Ref. [306] and Ref. [231] are two reviews of the field with similar scope published six years apart in the same journal, but the former cites ‘only’ 155 items, whereas the latter cites as many as 314 items. All this effort notwithstanding, the final word as to why cooperation evolves (among organisms in general and humans in particular) is yet to be uttered. Why is that?

Answering the posed question may not be straightforward, but some contributing factor can be singled out with confidence. Even after choosing

a social dilemma of interest, modelling studies need to make a series of assumptions about the modelled population and evolutionary dynamics in this population. Typical choices include finite or infinite [307, 308] and unstructured or structured [309, 310, 311] populations with fitness-based or pairwise-comparison [312, 313] dynamics. What is crucial here is that differing choices in a contextually similar situation can lead to widely different results [314, 315]. Models are, moreover, often proposed using cursory rationalisations that struggle to withstand the scrutiny of empirical tests [316, 317], which still somehow fails to discourage recycling similar rationalisations after the fact. Cases in point are peer punishment and network reciprocity. Peer punishment clearly promotes cooperation in theoretical models because getting punished is detrimental to payoff, but experimental results have been much less straightforward; despite an early confirmation [318], later research generated mixed [319] and negative [320] results, likely due to intimidatory and retaliatory uses of punishment. Similarly, network reciprocity has shown promising cooperation-promoting effects in theoretical models [242], particularly those built upon node-degree heterogeneous networks [250, 301], but experimental confirmations of such a promise had failed to materialise at first [321, 322] and remained constrained afterwards [323].

Theoretical and experimental advances in physics go hand in hand, and studying the evolution of cooperation should follow the same path. Why then have empirical tests so far failed to steer theoretical-model development? Perhaps this is in part because experiments involving human volunteers are far more ambiguous than experiments involving, say, elementary particles or gravitational waves. The above-mentioned examples of peer punishment and network reciprocity already show that behavioural experiments can go both ways. Ambiguity in empirical results is particularly unwelcome when central tenets of a theory are put to the test. One such tenet when considering the evolution of human behaviour is the imitation (i.e., update) rule [306]. Among early and popular imitation rules is the Fermi rule [324, 325] given by

$$p_{i \leftarrow j} = \frac{1}{1 + \exp\left(-\frac{\Pi_j - \Pi_i}{K}\right)}, \quad (54)$$

where $p_{i \leftarrow j}$ denotes the probability that the i th individual imitates the j th individual, with Π_i and Π_j being the respective payoffs of these individuals, and K denoting the irrationality of selection in the sense that when $K \rightarrow 0$,

then

$$p_{i \leftarrow j} = \begin{cases} 0, & \text{for } \Pi_i > \Pi_j \\ \frac{1}{2}, & \text{for } \Pi_i = \Pi_j, \\ 1, & \text{for } \Pi_i < \Pi_j \end{cases} \quad (55)$$

whereas when $K \rightarrow \infty$, then $p_{i \leftarrow j} = \frac{1}{2}$. Empirical evidence from economic game theory seems to provide strong support for imitation guided by payoff differences to the extent that volunteers consciously perceive themselves as imitators [326]. This being said, the nature and complexity of economic games precludes an immediate interpretation of actions by others as sympathetic or antagonistic to oneself (as is the case with cooperation or defection), which ultimately leaves little to go on apart from payoffs. In a comparative analysis of three spatial prisoner's dilemma experiments in which actions taken by others could readily be interpreted as sympathetic or antagonistic, payoff differences seem to have had no decision-making value [322]. Interestingly, one of the experiments analysed in Ref. [322] was used in a prior publication [327] to argue in favour of the Fermi rule. The latest work on the subject [328] is once again in favour of imitation because individuals confronted with more successful others imitate the behaviour of those others in accordance with the experienced payoff difference.

Just as it was the case with peer punishment and network reciprocity, empirical evidence favouring imitation driven by payoff differences comes with a degree of ambiguity. A plausible reason why this is so is that payoffs simply do not tell the whole story. For example, experiments reveal the existence of behavioural phenotypes [329], meaning that volunteer behaviours are not idiosyncratic, but rather exhibit recognisable characteristics. This enables classifying volunteers into a relatively small number of distinct groups or phenotypes. The cooperative phenotype, in particular, has been shown to possess remarkable robustness with respect to the form of cooperativeness and the passage of time [330]. Different people are, therefore, likely to be predisposed to cooperate to a different degree in the same social-dilemma situation. Volunteers may also directly respond to the actions of others instead of being concerned with payoffs. An antagonistic (respectively, sympathetic) action may provoke an antagonistic (resp., sympathetic) response. This is related to the theoretical concepts of the tit-for-tat strategy [331] and conditional cooperation [332], and indeed seems to regularly occur in behavioural experiments [333, 334, 335, 336]. Somewhat similarly, volunteers have been shown to cooperate more (respectively, less) after cooperating (resp., defect-

ing) previously even if the contextual situation is the same [335, 337, 338].

Additional reasons why payoffs may not tell the whole story in the context of cooperative behaviours is that human decision making is prone to some ‘peculiarities’. The field of behavioural economics made it a point to tease out such peculiarities, that is, document behaviours that deviate from the predictions of economic models based on a rational drive to maximise the expected utility [339, 340, 341]. Experiments focusing on cooperativeness also pinpoint peculiarities that are similar in spirit or related to those emphasised by behavioural economists. Ref. [342], for example, shows that acquaintances cooperate significantly more than strangers in a social-dilemma game in which there is absolutely no incentive to do so. Volunteers were, in fact, incentivised to score as much as possible for themselves following gameplay rules designed to be inconsequential to any sort of real-world interactions that may occur afterwards. The experiment was thus akin to point-gathering competitive gameplay that is so common among friends worldwide without ruining their friendships. That mere identifiability spurs altruistic behaviour has been observed elsewhere [343], but it is hard to incorporate into theoretical models when the corresponding incentives are entirely absent.

In another experiment, aimed at discerning how rewards promote cooperation, it was found that an unexpected and convoluted mechanism is at play [344]. Specifically, volunteers almost ignore the opportunity to reward one another, and yet cooperativeness doubles compared to the control in which there is no reward. To make matters even more perplexing, improved cooperation frequencies are observed right from the start, before any rewarding could ever happen. This peculiar behaviour can ultimately be traced to a known cognitive bias called the decoy effect [345]. Specific to multiple-choice situations, ‘decoy’ is a choice that shares some defining characteristics with another ‘target’ choice, but is inferior in one defining characteristic. Such inferiority makes the target look disproportionately more attractive, and thus preferred over all other choices. Because rewarding as envisioned in Ref. [344] was just an extra-demanding form of cooperating compared to the usual cooperative option, the decoy effect made cooperation appear much more attractive than defection, ultimately causing a surge in the cooperation frequency. On a more general note, cognitive biases may provoke not only behaviours for which there are no apparent incentives, but also behaviours that outright go against incentives implied by a particular social dilemma. This is particularly hard to model outside of broader evolutionary contexts in which cognitive biases may make more sense [346].

We have already mentioned in passing that from an empirical perspective, static networks offer a limited scope for promoting cooperation [323]. Consequently, the focus has shifted to dynamic networks in which volunteers could initiate links with cooperative others or sever links with defecting others [347]. Just as cooperation fails in many static networks, the same happens in networks that are shuffled randomly every round, or in networks that are rewired infrequently (albeit freely) by volunteers. Only when network links can be updated frequently and freely, a high degree of cooperation is maintained because cooperators shun links with defectors, while preferentially linking with other cooperators. This form of shunning defectors is often interpreted as a sort of punishment that, to make things even better, comes with two positive side effects. First, there is no obvious cost to severing links with defectors, which avoids what is known as the second-order social dilemma by which cooperators who refuse punishing are better off than cooperators who do punish [348, 349, 350]. Second, once the link is severed, there are no opportunities (at least not immediate ones) for retaliation, which avoids diminishing the willingness to punish defectors due to the threat of being retaliated against [351, 352, 353]. On a more fundamental level, however, dynamic networks can be seen as a means to foster positive assortment [354]. The importance of this perspective is that positive assortment has often been emphasised as a common denominator for many major cooperation-promoting mechanisms [236, 355, 356, 357, 358]. If positive assortment is indeed key to resolving social dilemmas, then enriching particularly hard social dilemmas with novel degrees of freedom that facilitate assortment should universally promote cooperation [336]. Thinking in terms of degrees of freedom is, of course, dear to physicists.

To summarise, empirical evidence shows that:

- In situations in which rationality dictates a clear course of action (e.g., behave to avoid punishment because it is bad for payoff), psychological factors may prompt another course of action (e.g., use punishment to intimidate or retaliate).
- The identification of distinct and stable behavioural phenotypes suggests that different individuals are predisposed to respond to the same social-dilemma situation with a different degree of cooperativeness.
- One's performance relative to those of others (in terms of payoffs) may be more influential in guiding decisions when the complexity of context-

tual situations conceals whether an action by another is sympathetic or antagonistic to oneself. When this is not the case, however, performance may be secondary to a direct response to the action by another.

- Intuitions may override incentives. For example, acquaintances are mutually more cooperative than strangers in exactly the same social-dilemma situation.
- Human decision making is fraught with peculiarities (i.e., cognitive biases) that make sense only in broader evolutionary contexts. These peculiarities may also (partly) override incentives. For example, a mere presence of reward promotes cooperativeness although no one rewards anyone.
- Introducing new degrees of freedom into hard social dilemmas may facilitate positive assortment and, by extension, promote cooperation. For example, static networks struggle to improve and maintain cooperativeness, but dynamic networks with frequent and free link updating work like a charm.

5.7. Future outlook

Having reviewed a wide variety of theoretical models pertaining to the evolution of human cooperation, and then summarising a number of empirical facts on the subject, we emphasised the need to reconnect theory and experiments. Doing so, however, faces challenges that have been first raised in relevant behavioural disciplines. Psychology and behavioural economics alike suffer from a deep replication crisis [359, 360, 361, 362, 363]. This state of affairs has triggered calls for an overhaul of the scientific process that had led to so many irreproducible results in the first place [364, 365, 366]. Many of the proposed measures are methodological; for example, the criterion for statistical significance should be more stringent [367], much larger samples are required to ensure high statistical power [368, 369], and preregistration should become a norm to curb statistical manipulations and avoid some cognitive biases that interfere with sound research practices [370, 371, 372]. All these measures are, if not necessary, then at least absolutely welcome, but they do complicate the logistics of conducting experimental studies (e.g., recruiting thousands of volunteers), and increase time and effort from conceptualising to publishing a study (e.g., preparing preregistration and calculating power). If physicists are to scrutinise their theoretical models by

means of behavioural experiments, the same high standards will apply as for psychologists, behavioural economists, and the like. Just keeping up with the standards will likely require widening multidisciplinary collaborations.

On top of methodological problems contributing to the replication crisis in behavioural disciplines, Ref. [373] makes a compelling argument that a lack of a strong theoretical backbone plagues the field even more. When such a backbone exists, it paints a bigger picture, guiding researchers to formulate useful expectations. Ref. [373] itself refers to an example from physics; early analyses of the data on neutrinos coming from CERN (near Geneva, Switzerland) to Gran Sasso National Laboratory (in the province of L'Aquila, Italy) suggested that these elementary particles move faster than light [374]. Because this would go against the special theory of relativity, the news of faster-than-light neutrinos was received with a healthy dose of scepticism. Later analyses indeed confirmed that neutrinos obey the limitation set by the speed of light [375]. It is almost a certainty that without firm theoretical expectations, the finding would be put to much less scrutiny, perhaps taking years to set the record straight.

Even more important is that, without firm theoretical expectations, an empirical behavioural scientist faces an almost infinite set of possible hypotheses that could be put to an experimental test. Experiments are then bound to be designed based on intuitions and guesswork about what may or may not be widely considered interesting (or worse yet, 'hot') among peers. A major consequence is that, even with impeccable experimental methodology, only bits and pieces of disconnected knowledge can be acquired. The research on the evolution of human cooperation is thankfully nowhere near such a dire state, but there are some elements reminiscent of what Ref. [373] refers to. Namely, evolutionary game theory specifies a blueprint for model construction that can always be extended with yet another cooperation-enhancing tweak. It is true that nuanced social-dilemma scenarios lead to fascinating dynamics that is of substantial interest in itself [262, 261], but do humans really behave as the models predict? Should we even empirically test countless modelled scenarios if they hardly get us any closer to general principles such as the aforementioned positive assortment? The answer would probably be negative even without the rising costs of behavioural experiments in terms of logistics, time, and effort. We therefore foresee the need for two types of models. One is models aimed at analysing incentives for *cooperative behaviour* in specific situations of societal interest such as fighting corruption [376, 377] or encouraging vaccination [378, 379]. The other is models aimed at explaining

the evolution of *cooperative trait* in humankind based on as general principles as possible such as robust paradigms for reciprocal altruism [331, 380] or robust social norms [381, 382]. Overall, this should lead to a narrower space of experimental hypotheses, theoretical models of cooperation that are backed by substantial empirical evidence, and ultimately more definitive answers as to when (circumstances) and why (forces) humans cooperate.

6. Networks and communities

Networks are a pillar of social physics. They permeate all aspects of the field, and more. Applications include—but are not limited to—online, physical, and even animal social networks [383, 384, 385, 386, 387, 388], finance [389], retail [390], supply chains [391, 392], transport infrastructure [393, 394, 395, 396, 397, 398], power grids [399, 400, 401], climate and Earth systems [402, 403], medical and clinical investigations [404, 405], nutrition [406, 407], and sports [408, 409]. Studies focusing on dynamics in networks are also ubiquitous, ranging from general dynamical patterns [35] to random walks [410, 411] to synchronisation [412, 413, 414], epidemiological dynamics [415, 416, 417], evolutionary dynamics of cooperation (see Section 5), social-balance dynamics [418, 419, 420], innovation dynamics [421], and many others. Network science applied to social systems has, in fact, grown to the point at which its branches are sufficiently broad to be a topic of massive standalone reviews [422, 423, 424, 425, 282, 426, 287, 427].

The sheer size of network science precludes us from overviewing the field comprehensively (let alone exhaustively). Our purpose is instead specific, that is, to examine the question of what constitutes a community in networks. We are particularly interested in decomposing and understanding the community structure of networks through the prism of generative models (as opposed to heuristic community-detection methods). For readers interested in the topic of community detection more broadly, modern developments and the current state of the art from the standpoint of physics can be found in Refs. [428, 423, 429, 430, 424, 431, 432]. The specific direction that we are singling out from the breadth of network science leads to at least two questions. What makes network-community structure so important? And why should one look favourably at generative models? We shall endeavour to answer these questions in the next section.

6.1. Community detection: contexts and methods

Detecting communities in networks is an intuitively appealing task. Thinking of a network as a means to consistently simplify the picture of a complex system, while retaining the ability to see how the system is stitched together into one whole, we may gain deeper insights into the functioning of the system by identifying the network's more basic constituents. Community detection can furthermore be considered a natural unsupervised-learning task (see Section 7.2). Indeed, when clustering is used to probe the internal structure of a dataset, among the first steps is to define a similarity or distance function between two data points (in essence forming a graph representation of the dataset), but networks already come with links that indicate relations between nodes. One underlying motivation behind community detection is therefore to leverage the preexisting information on network topology, and thus node-node relations, in order to learn about networks, much like an unsupervised learner learns about datasets. Given that communities as more basic constituents come together to form a network, it is of little surprise that the network's divisibility into communities shapes the dynamics that unfolds in networks. We have already mentioned this in the case of the evolutionary dynamics of cooperation [275, 276], but the same goes for, among others, epidemiological dynamics [433, 434] and related decision making [435, 436]. Arguably one of the most important features of community detection is the ability to make informed decisions about errors in measuring network structure. Doing so, however, demands having a 'standard' or a 'blueprint' that tells us whether we should expect a link where there is none or expect no links where there is one. Such a standard or blueprint is provided by generative models.

The importance of generative models is analogous to the importance of mechanistic (i.e., process-based) models in the context of dynamical systems. If we measure, say, the growth of a city, and the number of hospitals or schools necessary to sustain the city, we may gain the ability to plan for the future. If the number of hospitals increases sublinearly or linearly with city size, then the city growth is likely to be manageable. If, however, the number of schools increases supralinearly, then the city growth is at some point likely to deplete the resources needed for building and operating more schools. Although the ability to plan for the future is very much desirable, without a mechanistic model of city growth, we are in the dark as to what causes the number of hospitals to increase sublinearly (which is manageable) and the number of schools supralinearly (which is unmanageable).

Actionable insights may be gained by pinpointing the processes behind the sublinear increase in one case and the supralinear increase in the other. In a similar vein, a generative model that fits an empirical network very well may strongly favour the presence of a link when there is none in the data. Such a situation would give us much confidence that the link is missing due to a measurement error. We may even uncover the process that creates this particular link. Surely, the best known generative models in network science are Erdős-Rényi [437], Watts-Strogatz [438], and Barabási-Albert [439] (in which the processes of growth and preferential attachment decide network topology). When it comes to generating modular networks, stochastic block models have become a staple [440]. We shall rely heavily on this model type moving henceforward.

Before we introduce and define stochastic block models, and see how they are used for statistical inference, it is useful to look at four main community-detection contexts and method classes (Fig. 30). The reason for this is to show that the problem of community detection has no single correct formulation, let alone a single correct solution. For example, when designing a distributed computing system spread over several locations, the best partitioning is the one that minimises expensive long-distance links. Such a problem is best tackled using cut-based methods (Fig. 30A). If, by contrast, the aim is to understand the structure of large organisations based on social interactions, a collection of strongly interacting, and thus densely interconnected, individuals is likely to act as a functional group within the organisation. Clearly, this is a problem for clustering methods (Fig. 30B). When competing interests drive social interactions, strongly interacting, and thus densely interconnected, individuals may be opponents who belong to different teams. This is a type of problem for methods seeking stochastically equivalent nodes (Fig. 30C). Finally, if we aim to identify groups of individuals threatened by an epidemic, then interconnectedness comes secondary to epidemiological dynamics. Dynamic methods are expected to yield the best results (Fig. 30D).

The examples outlined above show that in community detection, context dictates methods. And yet, some methods are more heuristic and phenomenological, whereas others are more rigorous and fundamental. The latter is especially true of methods seeking stochastically equivalent nodes. These methods are founded on generative models and statistical inference. A prime example is stochastic block models to which we turn next.




Figure 30: Four main community-detection contexts and method classes. **A**, Cut-based methods seek to find network partitioning that minimises the number of between-community links without imposing dense within-community linking. **B**, Clustering methods embody the intuitive idea that links within communities are dense and across communities sparse. **C**, Methods seeking stochastically equivalent nodes posit that two nodes from the same community link to nodes from other communities with exactly the same probability. In the shown example, there are three communities identified by the block structure of the adjacency matrix. Community 1 has dense internal links, barely any links with community 2, and moderately dense links with community 3. **D**, Dynamic methods emphasise system behaviour over system topology. In particular, the role dynamics is considered crucial.

Source: Reprinted figure from Ref. [431] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

6.2. Introducing stochastic block models

As our discussion has already indicated, stochastic block models (abbreviated SBMs) power one of the most popular techniques for community detection that comes from the domain of statistical inference. The technique is based on the construction of an SBM that is fitted to network data [440, 441]. The model parameters are estimated by maximising likelihood, and once this is done, they provide information not only about the network structure, but also about the within-network node relationships, thus forming a flexible modelling tool for analysis and prediction. An additional advantage of SBMs over other methods is that SBMs are not limited to assortative mono-layered networks; it is conceptually straightforward to generalise the technique to a wide range of topological and dynamical network constructs (Fig. 31). Mixed memberships and overlapping communities [442, 443, 444], weighted networks [445, 446], multilayer networks [447, 448, 449, 450, 451], temporal networks [452, 453, 454, 455, 456, 457, 458, 459], networks that possess node attributes [460] or are annotated with metadata [461, 462] pose no problems to SBM-based community detection [463].

The first SBM algorithms were developed in social sciences to detect communities of ‘approximately equivalent’ nodes. The algorithms were de-

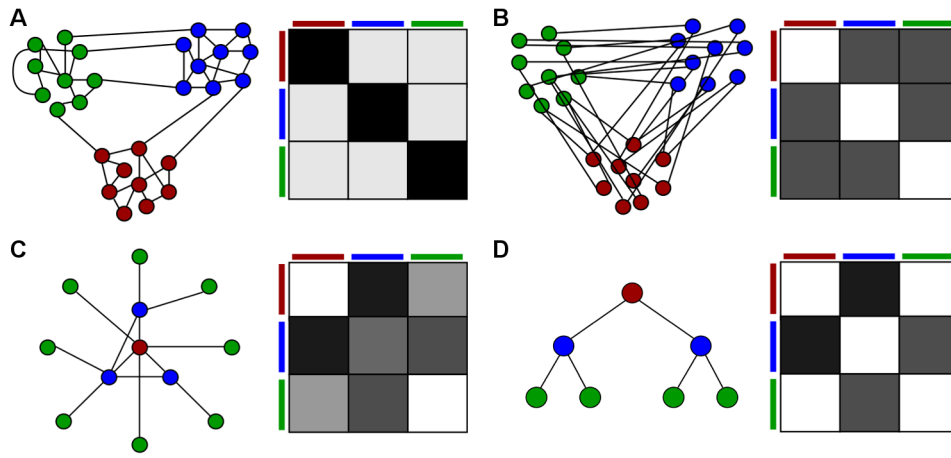


Figure 31: Representing various network topologies using stochastic block models. Each of the four networks has an adjacency matrix divided into blocks, where the grayscale indicates link probabilities (white=0 and black=1). **A**, Assortative network structure in which within-community links are dense, but between-community links are sparse. **B**, Disassortative network structure in which within-community links are sparse, but between-community links are dense. **C**, Core-periphery network structure. **D**, Hierarchical network structure.

Source: Reprinted figure from Ref. [464] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

terministic and were based on permuting the adjacency matrix to reveal the block structure and relationships among community members [465, 466, 432]. Shortly afterwards, pioneering works formalised the generative model and established a stochastic formulation of node equivalence in a community in such a way that equivalent nodes were associated with equivalent probabilities [440, 467, 441]. During the 1990s, computer power started to grow tremendously, which opened up the doors to new opportunities for network analysis in the digital environment. Consequently, SBM methods and algorithms underwent a rapid and diverse development over the last 25 years. SBMs are nowadays applied in a large number of natural and social scientific, and engineering fields (see Table I in Ref. [468]).

One of the most significant contributions to the development of community-detection methods based on statistical inference comes from physicists Karrer and Newman, who adapted the standard SBM model to take into account node-degree heterogeneity [469]. Thus obtained degree-corrected SBM enabled the application of SBMs to real-world networks and paved the way for

the development of multiple model variants [470, 471]. Newman, also the creator of one of the most popular heuristic methods based on modularity optimisation [472], in a recent study [473] showed that there is an equivalence between likelihood maximisation for SBMs and generalised modularity maximisation for planted-partition model with simplified community structure. The planted-partition model is a reduced version of a standard SBM in which links materialise with two probabilities, one for within communities and the other for between communities. Equivalence results such as Newman’s, and the extent to which they hold [474], are of much interest because they show that methods developed with very different motivations and intuitions in mind may end up serving the same purpose.

Another major contribution from statistical physics comes from Peixoto who developed a microcanonical view of SBMs by which the traditional (i.e., canonical), probabilistic definition of link formation between nodes belonging to two separate communities is replaced by a precise number of links [475, 463]. The canonical and microcanonical definitions are in accordance with the jargon of statistical physics; in the canonical generative model constraints on node degrees and the number of links are imposed on average, whereas in the microcanonical model these constraints are exact. One of the chief results in recent years is the development of a non-parametric microcanonical model using Bayesian inference that does not require prior knowledge of the number of communities, as well as the nested versions of this model for community detection [468, 463]. The result is a culmination of previous developments in the context of microcanonical SBMs, including the work on making use of minimum description length (MDL) [476, 477] to determine the number of communities [478], devising efficient Markov chain Monte Carlo inference algorithms [479], proposing nonparametric nested models [480], and incorporating model selection [481]. For a hands-on experience, most of the present SBM variants are implemented in *graph-tool*, a Python module for manipulation and statistical analysis of networks (available at <https://graph-tool.skewed.de/>).

In the rest of this introductory review, our focus will be on the above-mentioned concepts, whereas detailed reviews of the development of other SBM variants are given in Refs. [464, 482, 483]. Of note is that the development of new SBM variants is in many cases motivated by the specifics of real-world networks. The theoretical results related to establishment of the fundamental limits for community detection in the SBM, both with respect to information-theoretic and computational thresholds, are extensively

reviewed in Ref. [484].

6.3. Defining stochastic block models

The construction of SBMs as a generative network tool is based on the idea that network nodes are divided into communities, and that the existence of a link between two nodes is determined by the communities to which these nodes belong. These considerations impose conditions for generating ensembles that in a statistical sense represent a network. There are two main approaches to defining an SBM: (i) canonical, in which conditions are imposed on average, and (ii) microcanonical, in which conditions are imposed exactly.

Canonical form. The traditional formulation of SBMs is in canonical form. SBMs in canonical form are parameterised by two parameters, \mathbf{b} and \mathbf{W} , as follows; N nodes are distributed in K communities, and the affiliation of nodes to communities is expressed by the vector $\mathbf{b} = [b_i]$ of dimension N , where the value $b_i = r \in \{1, 2, \dots, K\}$ denotes the affiliation of the i th node to the r th community. The number of nodes in each community can be read from the vector \mathbf{b} . We denote the size of the r th community by n_r and form a vector of community sizes $\mathbf{n} = [n_r]$ of dimension K . The matrix $\mathbf{W} = [w_{rs}]$ of dimension $K \times K$ specifies the probabilities w_{rs} that a link is formed between any two nodes belonging to the communities r and s , that is, $P(i \leftrightarrow j) = w_{rs}$, where $b_i = r$ and $b_j = s$.

The matrix \mathbf{W} can be specified in multiple ways. When the values $w_{rs} \sim \mathcal{B}(p_{rs})$ follow a Bernoulli distribution with the parameter p_{rs} , then the probability that there is a link between nodes i and j is $P(i \leftrightarrow j) = p_{rs}$, and the probability that there is no link is $P(i \nleftrightarrow j) = 1 - p_{rs}$. In this case, the probability that a generative procedure, taking a node division \mathbf{b} as a parameter, creates an undirected and unweighted network with the adjacency matrix $\mathbf{A} = [A_{ij}]$ equals to

$$P(\mathbf{A}|\mathbf{p}, \mathbf{b}) = \prod_{i < j} p_{b_i, b_j}^{A_{ij}} (1 - p_{b_i, b_j})^{1 - A_{ij}}, \quad (56)$$

where $\mathbf{W} = \mathbf{p} = [p_{rs}]$ is matrix of dimension $K \times K$ comprising Bernoulli parameters.

A commonly used distribution for ease of calculation is the Poisson distribution. In this case, referred to as the standard SBM, the values $w_{rs} \sim \mathcal{P}(\lambda_{rs})$ are determined by the parameters λ_{rs} , which in turn define the probability

$P(i \Leftrightarrow j; k) = \lambda_{rs}^k e^{-\lambda_{rs}} / k!$ that there is a link of order k between nodes i and j . A network created using a Poisson distribution, unlike a Bernoulli one, has multiple links between nodes, which is convenient in the case of generating or analysing value networks. The Poisson distribution can, however, be used in conjunction with networks having just one link between two nodes. This is achieved by recognising that the probability of multiple links decreases with $1/N$ in sparse networks in which the total number of links is proportional to N , so the existence of multiple links can be ignored when N is large [463], or simply multiple links are compressed into a single one. The probability that the standard SBM generates a network with the adjacency matrix $\mathbf{A} = [A_{ij}]$ equals to

$$P(\mathbf{A}|\mathbf{\Lambda}, \mathbf{b}) = \prod_{i<j} \frac{\lambda_{b_i, b_j}^{A_{ij}}}{A_{ij}!} e^{-\lambda_{b_i, b_j}} \times \prod_i \frac{(\lambda_{b_i, b_i}/2)^{A_{ii}}}{(A_{ii}/2)!} e^{-\frac{\lambda_{b_i, b_i}}{2}}, \quad (57)$$

where $\mathbf{W} = \mathbf{\Lambda} = [\lambda_{rs}]$ is matrix of dimension $K \times K$ comprising Poisson parameters.

Degree-corrected SBM. The assumption behind the standard SBM is that nodes in the same community are statistically equivalent, that is, they all have the same number of links, on average. This is seldom true for real networks whose node-degree heterogeneity may span orders of magnitude [485]. To accommodate such node-degree heterogeneity, Ref. [469] proposes a modified model, called the degree-corrected SBM, in which each node is assigned the parameter θ_i controlling the node's expected degree $\langle k_i \rangle$, regardless of community affiliation. This means that in addition to the vector \mathbf{b} and the matrix $\mathbf{\Lambda}$, another model parameter $\boldsymbol{\theta} = [\theta_i]$, taking the form of an N -vector, is required to define the model. The parameter $\boldsymbol{\theta}$ generates heterogeneity within communities, where $P(i \Leftrightarrow j; k) = (\theta_i \theta_j \lambda_{rs})^k e^{-\theta_i \theta_j \lambda_{rs}} / k!$ is the probability that there is a link of multiplicity k between nodes $i \in b_i = r$ and $j \in b_j = s$. Given the additional parameter $\boldsymbol{\theta}$, a network with the adjacency matrix $\mathbf{A} = [A_{ij}]$ is generated with the probability

$$P(\mathbf{A}|\mathbf{\Lambda}, \boldsymbol{\theta}, \mathbf{b}) = \prod_{i<j} \frac{(\theta_i \theta_j \lambda_{b_i, b_j})^{A_{ij}}}{A_{ij}!} e^{-\theta_i \theta_j \lambda_{b_i, b_j}} \times \prod_i \frac{(\theta_i^2 \lambda_{b_i, b_i}/2)^{A_{ii}}}{(A_{ii}/2)!} e^{-\frac{\theta_i^2 \lambda_{b_i, b_i}}{2}}. \quad (58)$$

Microcanonical form. In canonical form, restrictions on node degrees and the number of links between communities r and s are ‘soft’, expressed via expected values. This means that across various model realisations, node degrees and the number of links fluctuate around the mean values. In micro-canonical form, the conditions are ‘hard’ in the sense that node degrees and the number of links are strictly determined for each realisation. Specifically, let the vector $\mathbf{k} = [k_i]$ of dimension N define node degrees, and furthermore let $\mathbf{e} = [e_{rs}]$ be a matrix of dimension $K \times K$ whose values e_{rs} define the number of links between communities r and s . For convenience, the diagonal elements e_{rr} are defined as double the number of links within community r . The generative process [486] assigns k_i semi-links (called ‘stubs’) to each node i , whereupon the stubs are randomly joined together until reaching the condition that between communities r and s there are exactly e_{rs} links. Connecting the stubs (i.e., wiring the network) to satisfy said condition can be done in $\Omega(\mathbf{e})$ ways

$$\Omega(\mathbf{e}) = \frac{\prod_r e_r!}{\prod_{r<s} e_{rs}! \prod_r e_{rr}!!}, \quad (59)$$

where $e_r = \sum_s e_{rs}$ and $(2m)!! = 2^m m!$. However, not every wiring produces a unique network. Given the adjacency matrix \mathbf{A} , the number of different stub wirings, $\Xi(\mathbf{A})$, that produce the same network is given by [486, 463]

$$\Xi(\mathbf{A}) = \frac{\prod_i k_i!}{\prod_{i<j} A_{ij}! \prod_i A_{ii}!!}. \quad (60)$$

This implies that the network with the adjacency matrix \mathbf{A} is generated with the probability

$$P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b}) = \frac{\Xi(\mathbf{A})}{\Omega(\mathbf{e})}. \quad (61)$$

The last relationship is valid under the ‘hard’ constraints

$$\begin{aligned} k_i &= \sum_j A_{ij}, \\ e_{rs} &= \sum_{ij} A_{ij} \delta_{b_i,r} \delta_{b_j,s}. \end{aligned} \quad (62)$$

If these constraints are not satisfied, then $P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b}) = 0$.

Relating canonical and microcanonical forms. As briefly mentioned before, equivalence results are important because they reveal how different motivations and intuitions may serve the same purpose. Even if full equivalence cannot be established, understanding conditions under which separate mathematical constructs exhibit similar behaviour is of great interest.

Canonical and microcanonical forms generate the same networks in the asymptotic sense if node degrees and the number of links between communities are large enough numbers. Namely, by expanding the relation in Eq. (58), we get

$$P(\mathbf{A}|\mathbf{\Lambda}, \boldsymbol{\theta}) = \prod_{r < s} \lambda_{rs}^{e_{rs}} e^{-\hat{\theta}_r \hat{\theta}_s \lambda_{rs}} \prod_r \lambda_{rr}^{\frac{e_{rr}}{2}} e^{-\frac{\hat{\theta}_r^2 \lambda_{rr}}{2}} \times \frac{\prod_i \theta_i^{k_i}}{\prod_{i < j} A_{ij}! \prod_i (A_{ii}/2)!}, \quad (63)$$

where $\hat{\theta}_r = \sum_i \theta_i \delta_{b_i, r}$. The parameters θ_i and λ_{rs} form a product in the expression for the probability $P(\mathbf{A}|\mathbf{\Lambda}, \boldsymbol{\theta})$, so their individual values can be re-scaled provided that the product remains the same. If we choose parameterisation

$$\hat{\theta}_r = \sum_i \theta_i \delta_{b_i, r} = 1 \quad (64)$$

for each community r , then $\lambda_{rs} = \langle e_{rs} \rangle$ is the expected number of links between communities r and s , and $\theta_i = \frac{\langle k_i \rangle}{\sum_s \lambda_{b_i, s}}$ is proportional to the expected node degree [463]. If furthermore the Stirling's factorial approximation $\ln(m!) \approx m \ln(m) - m$ is applied to Eqs. (59) and (60), it can be shown that the microcanonical likelihood $P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b})$ in Eq. (61) approaches asymptotically the likelihood $P(\mathbf{A}|\mathbf{\Lambda}, \boldsymbol{\theta})$ in Eq. (63), for large enough k_i and e_{rs} . For small or sparse networks, however, the differences between canonical and microcanonical forms can be substantial [487, 488].

As we have just seen, there is no exact equivalence between canonical and microcanonical forms. From the viewpoint of inference, however, such a lack of equivalence is immaterial because the models are unidentifiable anyway; it is impossible to tell from a single network realisation whether the network came from the canonical or the microcanonical model. Bayesian inference offers yet another perspective on the relationship between the two models, which will be discussed shortly.

6.4. Statistical inference of communities

Having defined the most common SBMs, it is now time to put them to good use, that is, use them for statistical inference. Given an observed network with the adjacency matrix $\mathbf{A} = [A_{ij}]$, $i, j \in \{1, 2, \dots, N\}$, statistical inference using SBMs consists of finding the model parameters that generate the observed network. More specifically, the problem is to find the node partition \mathbf{b} that maximises the log-likelihood function $\mathcal{L} = P(\mathbf{A}|\mathbf{b})$. Karrer and Newman [469] derived an unnormalised log-likelihood function for the standard model

$$\mathcal{L}^{\text{KN}} = \ln P(\mathbf{A}|\mathbf{b}) = \sum_{r,s} e_{rs} \ln \frac{e_{rs}}{n_r n_s}, \quad (65a)$$

where, as before, $e_{rs} = \sum_{ij} A_{ij} \delta_{b_i,r} \delta_{b_j,s}$ is the total number of links between communities r and s , or if $r = s$, double the number of links in community r . Peixoto [475] derived a similar result for the microcanonical SBM

$$\begin{aligned} \mathcal{L}^{\text{P}} &= \frac{1}{2} \sum_{r,s} n_r n_s H\left(\frac{e_{rs}}{n_r n_s}\right) \\ &\stackrel{\langle k \rangle \ll N}{\approx} -E + \frac{1}{2} \sum_{r,s} e_{rs} \ln \frac{e_{rs}}{n_r n_s}, \end{aligned} \quad (65b)$$

where $H(x) = -x \ln x - (1-x) \ln(1-x)$ is the binary entropy function. In the case of the degree-corrected SBM, the analogous log-likelihood relations for \mathbf{b} are (see Refs. [469, 475] for details)

$$\mathcal{L}_{\text{d.c.}}^{\text{KN}} = \sum_{r,s} e_{rs} \ln \frac{e_{rs}}{e_r e_s}, \quad (65c)$$

$$\mathcal{L}_{\text{d.c.}}^{\text{P}} = -E + \sum_k N_k \ln(k!) + \frac{1}{2} \sum_{r,s} e_{rs} \ln \frac{e_{rs}}{e_r e_s}, \quad (65d)$$

where $e_r = \sum_s e_{rs}$ is the total number of links with nodes affiliated with community r , and N_k is the number of nodes with the degree k . The likelihood functions in Eqs. (65a)–(65d), despite showing a way forward, suffer from a serious drawback. Naively minimising them would result in communities with only one node. It is therefore necessary to either know the number of communities K in advance or somehow estimate the value of K .

The methods for inferring the number of communities are diverse [482], but there are generally two dominant directions: (i) SBMs are fitted using

various values of the parameter K , and the optimal parameter value is determined by some measure or criterion; and (ii) K is not considered an external parameter, but is determined internally by the inference algorithm. The former approaches are called parametric, whereas the latter are called, somewhat anticlimactically, non-parametric. Parametric approaches oftentimes add to the likelihood function a part that acts as a penalty for the increasing number of communities. An example of this is the use of the Bayesian information criterion (BIC) [452, 489], the minimum-description-length (MDL) principle [478], or some other variation in terms of how to estimate likelihood or modify penalty [490, 491, 492].

Examples of integrating procedures for determining the number of communities K into inference algorithms are due to Newman and Reinert [493], who obtained a closed-form likelihood expression for the degree-corrected SBM, as well as Côme and Latouche [494] who used the exact integrated complete likelihood. Nowadays, however, among the most commonly used versions is the non-parametric microcanonical formulation due to Peixoto [468, 463].

Non-parametric microcanonical SBM. The total joint distribution for data and model parameters in microcanonical form is

$$P(\mathbf{A}, \mathbf{k}, \mathbf{e}, \mathbf{b}) = P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b})P(\mathbf{k}|\mathbf{e}, \mathbf{b})P(\mathbf{e}|\mathbf{b})P(\mathbf{b}), \quad (66)$$

where $P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b})$ is defined in Eq. (61), whereas $P(\mathbf{k}|\mathbf{e}, \mathbf{b})$, $P(\mathbf{e}|\mathbf{b})$, and $P(\mathbf{b})$ are prior distributions (Fig. 32). According to Bayes' theorem, the posterior distribution of a network partitioning into communities is

$$P(\mathbf{b}|\mathbf{A}) = \frac{P(\mathbf{A}|\mathbf{b})P(\mathbf{b})}{P(\mathbf{A})} = \frac{P(\mathbf{A}, \mathbf{b})}{P(\mathbf{A})} = \frac{P(\mathbf{A}, \mathbf{b})}{\sum_b P(\mathbf{A}, \mathbf{b})}, \quad (67)$$

where $P(\mathbf{A}, \mathbf{b})$ is the marginal joint distribution after integrating out the parameters \mathbf{k} and \mathbf{e}

$$P(\mathbf{A}, \mathbf{b}) = \sum_{\mathbf{k}, \mathbf{e}} P(\mathbf{A}, \mathbf{k}, \mathbf{e}, \mathbf{b}) = P(\mathbf{A}, \mathbf{k}', \mathbf{e}', \mathbf{b}), \quad (68)$$

where $\mathbf{k}' = \mathbf{k}'(\mathbf{A}, \mathbf{b})$ and $\mathbf{e}' = \mathbf{e}'(\mathbf{A}, \mathbf{b})$ are compliant with Eq. (62), that is, the hard constraints on node degrees and the number of links. Eqs. (66)–(68) offer conceptual guidance as to what we want to achieve through Bayesian statistical inference. Specifically, we want to maximise the probability $P(\mathbf{b}|\mathbf{A})$

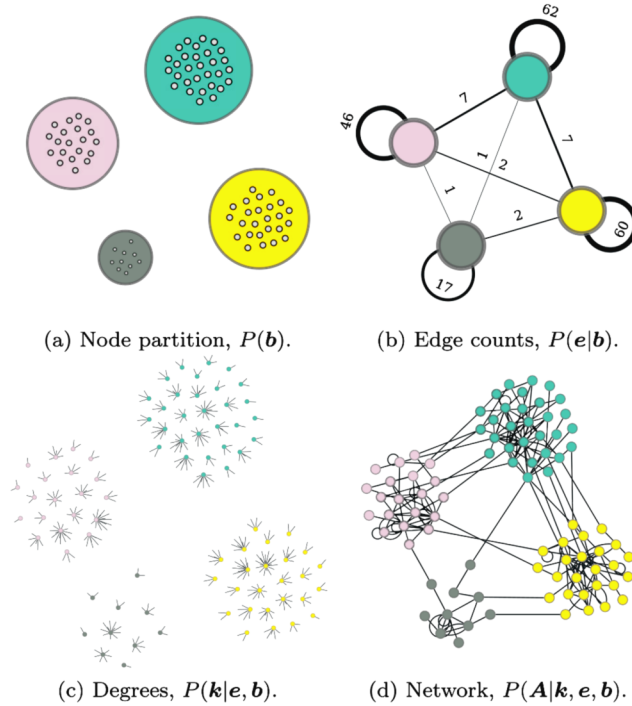


Figure 32: Non-parametric microcanonical stochastic block model and the corresponding generative process. Panel (a) illustrates the sampling of node partitions from the distribution $P(\mathbf{b})$. Panel (b) illustrates the sampling of edge counts from the distribution $P(\mathbf{e}|\mathbf{b})$. Panel (c) illustrates the sampling of node degrees from the distribution $P(\mathbf{k}|\mathbf{e}, \mathbf{b})$. Nodes are accompanied with semi-links or stubs that are yet to be wired into a network. Finally, panel (d) illustrates the sampling of the network from the distribution $P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b})$. Source: Reprinted figure from Ref. [468].

of partitioning \mathbf{b} conditional on observing the adjacency matrix \mathbf{A} . This turns out to be equivalent to maximising the joint distribution $P(\mathbf{A}, \mathbf{b})$ because the denominator in Eq. (67) acts just as a normalisation constant. The joint probability $P(\mathbf{A}, \mathbf{b})$, however, is knowable only via Eq. (66), implying that inference is a multi-part procedure. We need to specify all the prior distributions, and then maximise the resulting joint distribution.

The prior distributions are key ingredients of the inference procedure. Because in most cases there is no empirical information about priors, the prior selection is purposely kept uninformative. This prevents introducing bias to the posterior distribution, and allows the data to guide the partitioning of networks into communities. The following relations respectively define the

priors for the parameters \mathbf{b} and \mathbf{e} (see Refs. [468, 463] for details)

$$P(\mathbf{b}) = P(\mathbf{b}|\mathbf{n})P(\mathbf{n}|K)P(K) = \frac{\prod_r n_r!}{N!} \binom{N-1}{K-1}^{-1} \frac{1}{N}, \quad (69)$$

$$P(\mathbf{e}|\mathbf{b}) = \left(\binom{K(K+1)/2}{E} \right)^{-1}, \quad (70)$$

where $\binom{(\cdot)}{(\cdot)}$ is the multiset binomial coefficient, and E is the total number of links. The number of communities K is in this context called a hyperparameter. The corresponding distribution $P(K)$ is rather fittingly called a hyperprior. The prior distribution for the parameter \mathbf{k} is specified in one of two ways (see Ref. [468] for details)

$$P_u(\mathbf{k}|\mathbf{e}, \mathbf{b}) = \prod_r \left(\binom{n_r}{e_r} \right)^{-1}, \quad (71a)$$

$$P_h(\mathbf{k}|\mathbf{e}, \mathbf{b}) = \prod_r \frac{\prod_k N_k^r!}{n_r!} q(e_r, n_r)^{-1}, \quad (71b)$$

where N_k^r is the number of degree- k nodes in community r , and $q(m, n) = q(m, n-1) + q(m-n, n)$ with the boundary conditions $q(m, 1) = 1$ for $m > 1$ and $q(m, n) = 0$ for $m \leq 0$ or $n \leq 0$. Selecting the uniform prior $P_u(\mathbf{k}|\mathbf{e}, \mathbf{b})$ may cause that most nodes have similar degrees, in which case incorporating the heterogeneous prior $P_h(\mathbf{k}|\mathbf{e}, \mathbf{b})$ may help.

Bayesian equivalence. As already discussed, we fall short of achieving the exact equivalence between canonical and microcanonical forms, but the Bayesian framework offers another perspective on the subject. We start by marginalising the likelihood $P(\mathbf{A}|\mathbf{\Lambda}, \boldsymbol{\theta})$ from Eq. (63) as follows [468]

$$P(\mathbf{A}|\mathbf{b}) = \int P(\mathbf{A}|\mathbf{\Lambda}, \boldsymbol{\theta})P(\mathbf{\Lambda})P(\boldsymbol{\theta}|\mathbf{b})d\lambda d\boldsymbol{\theta}. \quad (72)$$

We have previously shown that, in the case of $\hat{\theta}_r = 1$ in Eq. (64), the values of λ_{rs} represent the expected number of nodes between communities r and s . If we choose a non-informative prior for λ in the form of an exponential distribution whose expectation is $\bar{\lambda} = 2E/K(K+1)$

$$P(\lambda_{rs}) = \begin{cases} \frac{1}{\lambda} e^{-\frac{\lambda_{rs}}{\lambda}}, & \text{if } r \neq s, \\ \frac{1}{2\lambda} e^{-\frac{\lambda_{rr}}{2\lambda}}, & \text{if } r = s, \end{cases} \quad (73)$$

and for $\boldsymbol{\theta}$ we choose another non-informative prior defined by

$$P(\boldsymbol{\theta}|\mathbf{b}) = \prod_r (n_r - 1)! \delta(\hat{\theta}_r - 1), \quad (74)$$

then from Eq. (72) it follows that [468]

$$\begin{aligned} P(\mathbf{A}|\mathbf{b}) &= \frac{\prod_{r<s} e_{rs}! \prod_r e_{rr}!!}{\prod_{i<j} A_{ij}! \prod_i A_{ii}!!} \\ &\times \prod_r \frac{(n_r - 1)!}{(e_r + n_r - 1)!} \prod_i k_i! \\ &\times \frac{\bar{\lambda}^E}{(\bar{\lambda} + 1)^{E+K(K+1)/2}} \\ &= P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b}) P(\mathbf{k}|\mathbf{e}, \mathbf{b}) P(\mathbf{e}). \end{aligned} \quad (75)$$

The quantity $P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b})$ is precisely the microcanonical likelihood from Eq. (61), the quantity $P(\mathbf{k}|\mathbf{e}, \mathbf{b})$ is precisely the uniform prior from Eq. (71a), whereas the quantity

$$P(\mathbf{e}) = \frac{\bar{\lambda}^E}{(\bar{\lambda} + 1)^{E+K(K+1)/2}} \quad (76)$$

differs from the microcanonical prior in Eq. (70) only in the sense that the total number of links may fluctuate, but in expectation the canonical and microcanonical forms are equivalent to one another [468].

Minimum-description-length interpretation. The non-parametric microcanonical SBM can be reinterpreted from an information-theoretical perspective, which offers an intuitive explanation of why this model is robust to overfitting. Namely, we can write $P(\mathbf{A}, \mathbf{k}, \mathbf{e}, \mathbf{b}) = 2^{-\Sigma}$, where by taking the logarithm of both sides, we get a value called the description length of data [476]

$$\begin{aligned} \Sigma &= -\log_2 P(\mathbf{A}, \mathbf{k}, \mathbf{e}, \mathbf{b}) \\ &= -\log_2 (P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b}) P(\mathbf{k}, \mathbf{e}, \mathbf{b})) \\ &= \mathcal{S} + \mathcal{L}. \end{aligned} \quad (77)$$

The quantity $\mathcal{S} = -\log_2 P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b})$ equals the number of bits needed to describe a network when the model parameters are known, and the quantity $\mathcal{L} = -\log_2 P(\mathbf{k}, \mathbf{e}, \mathbf{b})$ equals the number of bits needed to describe the

model. By maximising the joint probability distribution in Eq. (66), a set of parameters is automatically obtained that gives the minimum description length. Moreover, the increasing value of \mathcal{L} acts as a penalty on the number of parameters, thus also limiting the number of communities. Without such a penalty the increasing number of communities tends to decrease the value of \mathcal{S} , meaning that ultimately each node would comprise its own community. In sum, the MDL interpretation shows that the non-parametric microcanonical SBM is a formal implementation of Occam’s razor, according to which the simplest model with a sufficient significance level is to be preferred.

Resolution limit and the nested SBM. Community detection by means of SBMs comes with a resolution limit, meaning that communities below the resolution minimum will not be assigned sufficient statistical significance. Instead, such communities will be merged into larger ones. A well-known example is a network of 64 10-node cliques (i.e., complete subgraphs) in which the cliques are mutually disconnected (i.e., there’s no links between any two cliques). Fitting the microcanonical SBM yields 32 communities, each comprising two cliques [463]. The model thus suffers from underfitting that manifests as an inability to detect communities below the resolution limit, which in turn scales with $\mathcal{O}(\sqrt{N})$ [478].

Peixoto [480] offered a solution to the resolution-limit problem in the form of a nested SBM. This type of SBM is based on a simple idea that the communities and the number of links between them, as determined by fitting an SBM, form a new multigraph (i.e., a network in which any two nodes may be linked multiple times, including closed loops). According to this idea, the communities represent the nodes of the new multigraph, the number of links between any two communities represents link multiplicity, and the number of links within a community represents loop multiplicity (Fig. 33). It is then possible to fit the SBM to the new multigraph again, producing yet another multigraph. By repeating the procedure recursively, we get a smaller and smaller number of communities until we finally reach the multigraph with one community. The reason why this method improves the resolution limit is that a higher-level multigraph serves as the information prior for the next lower level. The method generalises the ‘flat’ model described above, and is applicable to large networks, be they assortative or disassortative.

Inference algorithms. Algorithms that effectively infer community affiliations often come from the Monte-Carlo class of methods used in statistical physics.

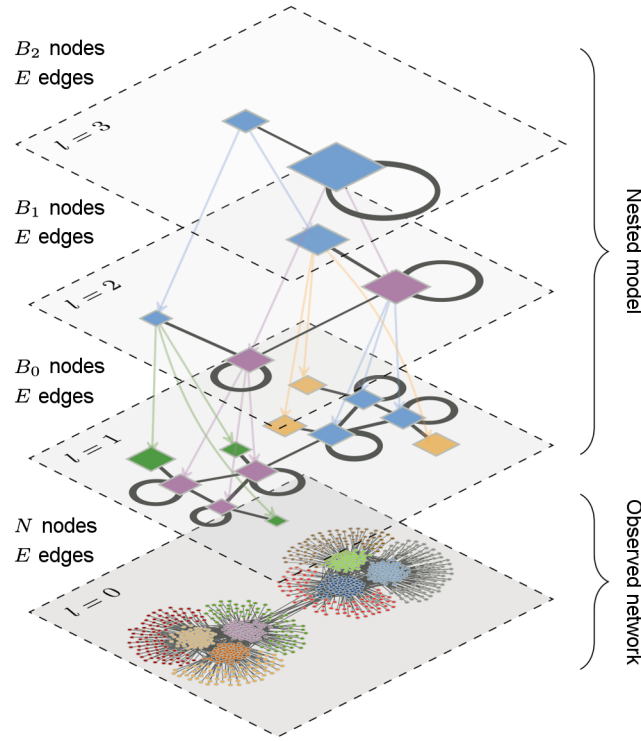


Figure 33: Nested stochastic block model in action. Stacked on top of the observed network are the corresponding three levels of multigraph representation ($l = 1$ to $l = 3$). To break the resolution limit, a higher level serves as a prior for the next lower level. Source: Reprinted figure from Ref. [480] under the Creative Commons Attribution 3.0 Unported (CC BY 3.0).

Although for different SBM variants exact expressions for the posterior probability can be derived in accordance with Eq. (67), up to the normalisation constant, the distributions are in most cases quite complicated. Therefore, Markov-chain Monte-Carlo (MCMC) methods for sampling from complicated distributions are proving to be an efficient and easy-to-implement inferential tool.

Examples of MCMC algorithms are Metropolis or Metropolis-Hasting algorithms [495, 496]. In the Metropolis algorithm, to sample from our target distribution $P(\mathbf{b}|\mathbf{A})$, we initialise the algorithm at an arbitrary position $\mathbf{b} = \mathbf{b}_0$. Next, a candidate replacement \mathbf{b}' is sampled from a symmetric, but otherwise arbitrary, distribution; if this distribution is normal and centralised around \mathbf{b} , choosing the candidate replacement \mathbf{b}' amounts to mak-

ing a random-walk step. We then calculate the ratio $f(\mathbf{b}')/f(\mathbf{b})$, where the function f needs only to be proportional to our target distribution $P(\mathbf{b}|\mathbf{A})$. This proportionality requirement is, in fact, one of the main strengths of the Metropolis algorithm because calculating the normalisation constant for the target distribution is often a non-trivial task. If said ratio is greater-or-equal than unity, i.e., $f(\mathbf{b}')/f(\mathbf{b}) \geq 1$, then the replacement candidate \mathbf{b}' comes from a region that is more probable under $P(\mathbf{b}|\mathbf{A})$ than the region around \mathbf{b} ; the candidate replacement becomes the new current sample. If, by contrast, the ratio is lower than unity, i.e., $f(\mathbf{b}')/f(\mathbf{b}) < 1$, then the replacement candidate can still be accepted as the new current sample, but the probability of doing so decreases as the ratio gets closer to zero. Occasional acceptance of candidate replacements that are less probable under $P(\mathbf{b}|\mathbf{A})$ is necessary because ‘less probable’ is still possible. Once a sufficiently large sample is obtained, the algorithm can be stopped, but of note is that there is no natural termination criterion. The Metropolis-Hastings algorithm relaxes the condition that candidate replacements need to be drawn from a symmetric distribution. Beside MCMC algorithms, frequently used are variational [497, 498, 499, 500] and greedy methods [469, 494, 490], which have been extensively reviewed in Ref. [482].

6.5. Future outlook

Community detection is a fast-paced research domain whose explosive development over the past 20 years or so stems from mathematical foundations that were laid decades ago. We started our overview of research on community detection with four main community-detection contexts, meanwhile emphasising the importance of generative models for consistent statistical inference of communities from network data. We singled out SBMs as a poster child for tremendous progress that had been made. Here, we recognise that as the SBM-based theory has been maturing, new methodological advances started taking roots and shaping the field’s future [483].

Much effort has been put into reformulating the problem of community detection to conform to the format of some machine-learning technique. Examples of this are community detection using topic models or matrix factorisation [483]. Topic models originate in machine learning and natural-language processing, and are based on an idea that documents, as word collections, refer to a limited number of topics. A topic in this approach is a cluster of similar words. The main task of a specific topic model is to map a set of documents into word-use statistics, and from there make two

inferences. One inference is the collection of topics that the documents collectively cover. The other inference is how much a certain topic is represented in a given document. Community detection based on topic modelling thus implies packaging nodes and links as words and documents for the model to process. The returned topics are then interpreted as communities. In a similar fashion, matrix factorisation is a class of collaborative-filtering algorithms devised for recommender systems with the goal to learn low-dimensional representations of users and items that can be used to predict how users rate items (e.g., how subscribers rate shows on a streaming service). Community detection based on matrix factorisation treats the adjacency matrix as ratings, while low-dimensional representations to be learned are those of links outgoing from and incoming into a fixed number of communities. In practice this means that an $N \times N$ adjacency matrix is decomposed into a product of an $N \times K$ matrix of outgoing links and $K \times N$ matrix of incoming links. The latter two matrices respectively give probabilities that the i th node generates an outgoing link from community r and that the j th node receives an incoming link into community r , which ultimately determines the community structure.

With the recent rise in popularity of neural networks and deep learning (see Section 7.3), it is unsurprising that the problem of community detection has also been cast in the form suitable for deep neural networks [483]. The main idea in this context is to learn node representation, that is, extract important features that set nodes apart, including their community affiliations. In deep neural networks, features are encoded by hidden network layers, but the whole business of feature extraction is perhaps easier to present by referring to a more ‘manual’ approach in which network nodes are embedded into a low-dimensional vector space. Such an embedding is achieved by first defining a node-similarity measure on networks. A well-known example is unbiased, fixed-length random walks that measure similarity in terms of the probability of visiting the node v during a random walk starting from the node u [501]. Second, a map (i.e., an embedding) is defined that associates network nodes with vectors in the vector space. Lastly, the parameters of this embedding are optimised in such a way that the similarity measure on networks is well approximated by (some function of) the scalar product in the vector space. This exact problem has, in fact, been rigorously treated by mathematicians with the diffusion distance being the similarity measure defined on networks, diffusion maps serving as the embedding, and the Euclidean distance approximating the diffusion distance in the low-dimensional

vector space [502]. In practice, however, measuring node similarity is expensive and therefore done locally. The embedding is optimised using only measured (as opposed to all) node similarities. Once this is achieved, community detection takes place by clustering vectors, corresponding to network nodes, in the low-dimensional vector space. Ref. [503] demonstrates the approach in action.

Among the briefly described approaches to community detection inspired by machine learning, deep neural networks in particular shed the connection to generative processes and statistical inference. This, however, is unlikely to slow down further proliferation of such methods. If anything, the most practical methods in terms of the ability to handle large network sizes (e.g., billions of nodes) and various network types (e.g., multilayer, dynamic, and incomplete) are likely to prosper in the future. We nonetheless expect community-detection methods founded on generative processes and statistical inference to continue going strong due to fundamental advantages and unrivalled rigour.

7. Human-machine networks

Human society is currently experiencing the impact of a digital transition by which data about human behaviour has evolved from a limited and unused resource to a manifold of permanently growing real-time data streams called *big data*. Today, big data is being pervasively generated, collected, analysed, and utilised within various smart systems to support enjoyable and comfortable living and working conditions. We are, in fact, witnessing a transition from *information* to *knowledge society* that has prompted the emergence of technologies whose potential to unravel both individual and collective behavioural phenomena is unprecedented. Examples of such technologies include:

- Google Knowledge Graph augmenting web navigation of Internet users [504],
- Facebook Social Graph revealing users' personal relations [505],
- LinkedIn Economic Graph digitally mapping every member of the workforce [506],
- Unacast Real World Graph explaining how people move around the planet [507], and

- Pinterest Taste Graph visually exploring what people like and what inspires them [508].

Within the ongoing digital transformation, the ambition of ensuring smooth progress along the data-information-knowledge-wisdom hierarchy [509] undoubtedly stays one of the major challenges given the ever-growing size, diversity, and frequency of generated data. The vehicle for navigating such a knowledge hierarchy is *Data Science*, an interdisciplinary field dealing with various methodologies and technologies (i.e., algorithms and systems) to automatically derive knowledge from data. Understanding and augmenting human intelligence and human-decision making processes is the necessary next step to harness data science in the pursuit of *actionable knowledge* for improving human lives.

Due to their influence on human lives, the evolution of data-science methodologies and technologies is intertwined with social and organisational structures in what is known as *socio-technical systems*. From this viewpoint, novel technologies emerge to satisfy societal needs, all the while society adapts to better accommodate those technologies (Fig. 34). For example, Internet as a global packet data network and the World Wide Web as an information system provide critical services for modern-day society, yet neither were originally designed with such broad purposes in mind.

Some of the emerging and evolving technologies within the framework of socio-technical systems are Distributed Ledger Technology (e.g., blockchain) [510, 511], Cyber-Physical Systems [512], Internet of Things [513], Cloud, and post-Cloud paradigms Fog, Edge, and Dew Computing [514]. The post-Cloud paradigms in particular strive to relocate computing resources closer to end users to mitigate cloud-related issues of highly centralised computation [515]. All these technologies, together with advances in the research of Artificial Intelligence (AI) [516, 517, 518] already influence everyday lives [519] to the point that humans and machines are entangled into elaborate *human-machine networks* [520]. Because of their increasing complexity and relevance in modern societies, human-machine networks are among the most challenging, and yet, most important environments to study human and machine co-behaviour [521].

The need to understand human behaviour in conjunction with big data collected from a variety of human-machine networks has given rise to a new discipline called *Computational Social Science* [522, 523]. This discipline encompasses modern trends in social-physics research [524, 525] based on the




Figure 34: Conceptual illustration of the evolution of socio-technical systems (STS) and human-machine networks (HMN) driven by emerging technologies. Although the inter-connection of structures and functionalities ensures scalable development, there is also the risk of a single local failure in one subsystem provoking a cascade of failures throughout other subsystems. For example, a failure of the power grid can cause the failure of information and communication technology systems, upon which financial, healthcare, or security services may depend down the line. Further vulnerabilities include the spread of disinformation or cyber-attacks. Ensuring not only scalability, but also robustness and security of human-machine networks therefore constitutes a critical task within the framework of social-technical systems.

joint use of computational big-data analyses on the one hand, and models from physical sciences on the other hand. Specifically, methods are borrowed from behavioural economics and social psychology, network science, data science, and machine learning, as well as game theory and the theory of critical phenomena. Discoveries are thus possible on three distinct levels:

1. *Data analyses* generate insights directly from collected data,

2. *Modelling* attempts to capture plausible governing mechanisms and processes, and
3. *Simulations* yield system-wide or component-wise predictions of behaviour within the human-machine network of interest.

The goal is to develop expressive-yet-simple models that can be calibrated and validated against real-world data as opposed to using phenomenological models that are limited to generic insights into governing mechanisms and processes. This type of data-driven modelling overcomes the drawbacks of black-box machine-learning algorithms in which the underlying physical principles of social systems remain entirely neglected.

Data-driven modelling as described above opens the door to deep insights into human behavioural patterns and better decisions in response to critical social problems. Areas of potential betterment include monitoring socio-economic deprivation of individuals and countries [526], increasing public wealth and health [527, 528], controlling safety and crime [529], mapping epidemics [530, 531], managing natural disasters [532, 533], and securing social inclusion [534, 535]. Multiple blooming research directions have caused much excitement, culminating in the idea of *social-good algorithms* [536, 537, 538, 539] that should guide all aspects of sustainable development [539], decision-making, and resource optimisation [537, 538].

Due to their supposed influence over so many aspects of modern-day life, social-good algorithms are undoubtedly powerful. Yet, with great power comes great responsibility. Concerns have already been raised around a range of social, ethical, and legal issues, including privacy and security [540, 541], transparency and accountability [542], and discrimination and bias [543]. The fact that people remain largely unaware of how algorithms utilise their data and affect their lives has been encapsulated in the expression *black-box society* [542]. Especially today's Internet-enabled human-machine networks such as online social networks, search engines, or other cloud-based platforms are highly centralised, exposing users' personal data to potential commercial or political misuses. Adding to the mix AI systems built upon complex deep-learning models, infamous for their lack of transparency, accentuates the need to tread carefully in the near future.

A key challenge for researchers and policymakers, in order to avoid the pitfalls of a black-box society, is to ensure maximum transparency in ever-growing interactions between humans and machines. Attempts to rise to this challenge have culminated in the framework of *Trustworthy Artificial*

Intelligence, at the core of which lies recognition that AI systems must be more easily interpretable in general, and explainable to various user groups in particular [544]. Further requirements in this context are alignment with fundamental human rights and legal practices, as well as service to societal common good [545, 546]. The European Union (EU) through the European Commission’s High-Level Expert Group on AI has taken an early lead in setting the pathway towards trustworthy AI; for example, the document “Ethics Guidelines for Trustworthy AI” [547], published in April 2019, provides concrete guidance on how to operationalise the above-stated requirements in future human-machine networks.

Trustworthy AI, with its three key characteristics (lawful, ethical, and robust) and seven key requirements (human agency and oversight, technical robustness and safety, privacy and data governance, transparency, diversity, non-discrimination and fairness, environmental and societal well-being and accountability) clearly sets the research-and-development direction for emerging AI-driven human-machine networks. Especially the focus on human-understandable algorithms that manipulate various types of unstructured, semi-structured, or structured data is sowing the seeds of an indispensable tool in the toolbox [548] of computational branches of social physics. These promising developments and a growing interest in AI research notwithstanding, existing studies about collective human [525] and machine behaviour [521], and their symbiotic interdependence [520, 549, 550], are highly fragmented. The fragments, comprising Human-Imitative AI, Intelligence Augmentation, and Intelligent Infrastructure, are complementary to one another and should fuse together in future studies of large-scale human behaviour.

The aim of this chapter is to help aspiring social physicists (i) to navigate, at times chaotic, advances in the domain of AI-driven human-machine networks, and (ii) to identify research directions where future breakthroughs may lie. To this end, the following sections start with an up-to-date review of studies at the interface between human [525] and machine [521] behaviour. Thereafter, the methodological fundamentals of AI are laid out in a form condensed for easy understanding. After a brief mention of exemplary uses in social-good and sustainable-development contexts, the AI methodology is exemplified in detail via the use of AI agents in game theory, and especially for the purpose of promoting cooperation. The chapter concludes with an outlook for the future.

7.1. Literature walkthrough

The use of machine learning in general, and deep learning in particular, to understand large-scale human behavioural phenomena has traditionally been scattered among research communities. Human and machine behaviour have thus been researched independently for the most part. Only recently an interface between these research topics has started to emerge. To gradually zero in on this interface, we categorise the literature into (i) general AI, machine-learning, or deep-learning techniques, (ii) AI, machine-learning, or deep-learning techniques to address societal challenges, (iii) AI, machine-learning, or deep-learning techniques in social physics, and (iv) collective human-machine behaviour.

There exists a large number of machine-learning and deep-learning surveys tailored to the needs of both specific [516, 518, 551, 552, 553, 554] or general audiences [555, 556, 557, 558, 559, 560, 561, 562]. Prioritising the latter, Ref. [558] presents an overview of most popular models and provides a long-term outlook for the field. Ref. [559] offers a comprehensive historical overview of the relevant work that deep learning builds on. Ref. [557] is notable for presenting deep-learning applications to a variety of information processing tasks. The topic of AI for social good, aiming at advancing and employing AI, machine learning, or deep learning to address societal challenges, is covered in Ref. [538]. Considering the ongoing Covid-19 pandemic, Ref. [563] overviews recent studies that utilise machine learning to tackle aspects of the pandemic at different scales, including molecular, clinical, and societal. Yet other recent studies have investigated economic [564, 565], social [566], as well as psychological and mental impacts [567] of the drastic life changes brought about by the pandemic. Lastly, Ref. [552] surveys recent developments in deep learning for recommender systems, which are currently one of the most established AI, machine-learning, or deep-learning applications within human-machine networks, having an important role in many online services and mobile apps.

Depending on whether machine learning is conducted with or without labelled input-output example pairs, we respectively distinguish between supervised and unsupervised learning. An intermediate approach, called semi-supervised learning, is useful in applications when unlabelled data is readily available or easy to acquire, while labelled data is often expensive or otherwise difficult to collect [568]. Ref. [569] is a comprehensive overview of recent advances in the domain of semi-supervised deep-learning techniques. Meanwhile, in the domain of unsupervised deep learning, progress has ma-

terialised in the form of generative models such as Variational Autoencoders and Generative Adversarial Networks. The use of the former in deep-learning contexts is covered in Ref. [570], while the latter, with applications, is covered in Ref. [571].

Deep reinforcement learning is a core AI research direction aimed at solving complex sequential decision-making tasks with potentially wide-ranging applications such as robotics, smart infrastructure, healthcare, finance, and others. Ref. [572] offers an introduction to deep reinforcement-learning techniques, models, and applications, while Refs. [573, 574] represent more comprehensive guides into the field. Furthermore, there exist numerous resources that allow for a more hands-on approach:

- “Spinning Up in Deep RL” is a practical introduction [575],
- OpenAI Gym is a collection of benchmark problems (i.e., environments) for comparing deep reinforcement-learning algorithms [576],
- OpenAI Baselines is a set of baseline implementations of deep reinforcement-learning algorithms [577], and
- rlpyt is an open-source repository of modular and parallelised implementations of various deep reinforcement-learning algorithms [578].

More recently, research on deep reinforcement learning has taken a turn from single-agent to multi-agent scenarios [579, 580]. Four topics of interest have crystallised in this context: (i) *the analysis of emergent behaviours* pertains to evaluating single-agent deep reinforcement-learning algorithms in multi-agent scenarios, for example, cooperative, competitive, and mixed; (ii) *learning communication* pertains to agents learning both through actions and messages; (iii) *learning cooperation* pertains to agents learning to cooperate using only actions and (local) observations; and (iv) *agents modelling agents* pertains to agents reasoning about other agents to fulfil a task, for example, cooperative or competitive. Cooperative tasks between actors in human-machine networks are of particular interest to the social-physics agenda and will, therefore, be discussed in more detail later.

Turning to the specifics of the AI, machine-learning, or deep-learning use in science, Ref. [553] overviews the techniques for applying deep-learning models in conjunction with limited data (self-supervision, semi-supervised

learning, and data augmentation), as well as the techniques for interpretability and representation analyses. Ref. [581] discusses the natural-science applications of explainable machine learning via three core concepts: *transparency*, *interpretability*, and *explainability*. Finally, Ref. [544] is a guide into the explainable deep learning aimed at researchers just entering the field.

Alongside many other branches of science, physics has not been immune to adopting the machine-learning methodology [548, 582]. Statistical physics has, in fact, inspired the exploration and development of machine-learning models [583, 584, 585, 586, 587]. A recent comprehensive review introduces the key concepts and tools of machine learning in a physicist-friendly manner accompanied with a set of Python Jupyter notebooks that demonstrate the application of modern machine-learning and deep-learning packages on physics-inspired datasets [554]. Refs. [588, 589] cover, from a theoretical perspective, the intersection between the foundational machine-learning and deep-learning concepts and statistical mechanics. The methods from statistical mechanics have furthermore begun to provide conceptual insights into deep-learning models regarding the model expressivity [590, 591, 592, 593], the shape of the model loss landscape [594], model training and information propagation dynamics [595, 596, 597], model generalisation capabilities [598, 599], and the model ability to ‘imagine’, that is, build deep generative models of data [570, 600].

To conclude this literature walkthrough, there is an immensely rich body of literature on the AI, machine-learning, and deep-learning techniques, and the contribution of physicists to this richness has been non-negligible to say the least. But where should an aspiring social physicist look for potential breakthroughs? Ref. [521] proposes a new field of scientific study called *machine behaviour* to better understand how AI agents might affect society, culture, the economy, and politics. The four primary motivations for the study of machine behaviour are the ever-increasing ubiquity of AI algorithms in human daily activities, but also, their complexity, opacity, and a lack of explainability. This ‘black-box’ nature of AI algorithms poses substantial challenges to predicting the effects of such algorithms, whether positive or negative, on humanity and society. There are three scales at which to study machine behaviour: individual machines, machine networks, and human-machine networks. Ref. [520] surveys the state-of-the-art developments on the third scale, identifying eight different types of human-machine networks depending on structure and interactions. These eight types are public-resource computing, crowdsourcing, web-search engines, crowdsensing, online markets, social me-

dia, multiplayer online games and virtual worlds, and mass collaborations. Nowadays, however, the omnipresence and usability of human-machine networks is causing novel trends to emerge by which the limits between the eight listed types are beginning to blur, while hybrid types keep cropping up. This state of affairs suggests a long and winding road, and thereby plenty of opportunities, towards the ultimate goal of the aforementioned social-good algorithms.

7.2. *Fundamentals of artificial intelligence*

Here, we introduce some of the most fundamental terms and concepts of AI, machine learning, and deep learning. Readers who wish a first-hand experience with the methods and techniques that arise from these concepts may want to consult Ref. [601] as a form of a getting-started tutorial.

Historically, the term ‘artificial intelligence’ was introduced in the late 1950s to refer to the aim of making intelligent machines that have the high-level cognitive capability to act, reason, think, and learn like humans, that is, to mimic human intelligence [602, 560, 603]. The evolution towards this aim of *human-imitative AI* further led to the emergence of AI systems that *augment human intelligence*, as well as those that constitute *intelligent infrastructure* in order to make living and working environments safer and more supportive of human needs [603]. Distilled to these three complementary research activities, that is, human-imitative AI, intelligence augmentation, and intelligent infrastructure, the research on AI comprises a very broad and diverse set of techniques for building and integrating intelligent agents into software solutions and hardware platforms.

Because human intelligence is general, the aim of achieving complete human-imitative AI is often called artificial general intelligence or the ‘strong’ AI [604]. Achieving general intelligence that would encompass all, or most, of human cognitive processes is beyond the reach of current AI research [605], meaning that presently only the ‘weak’ AI systems are realisable. These weak AIs are also referred to as artificial narrow intelligence. A further implication is that machines learn to perform good on a specific, well-defined task, but cannot augment humans outside of a limited domain in which the machine learned to operate. Limitations notwithstanding, modern-day artificial narrow intelligence is widely used in various domains, ranging from science to business to health care and more. Interestingly, most systems that classify as AI today are, in fact, based on the machine-learning and deep-learning techniques (Fig. 35).




Figure 35: A Venn diagram showing overview and composition of AI technology. Deep learning is a part of machine learning, which in turn is a part of AI. Machine learning is commonly divided into supervised learning, unsupervised learning, and reinforcement learning. Deep learning for the most part inherits this division. The major benefit of deep learning over traditional machine learning is the automatic feature extraction that circumvents expensive feature engineering by hand.

Machine learning as a subfield of AI enables artificial agents (i.e., machines) to automatically learn from data, make decisions and predictions by themselves, and help in human decision making without being explicitly programmed with expert knowledge [606, 554]. This ability of AI systems to automatically extract new knowledge from data is an extension of knowledge-based approach to AI by which knowledge about the world is to be hard-coded using formal languages, while machine reasoning should follow logical inference rules on statements formulated in such languages. In practice, a successful machine-learning algorithm recognises important features in a ‘training’ dataset in order to make inductive inferences or predictions about data samples unseen during training. The machine-learning algorithm must thus ‘generalise’ beyond data in the training dataset; the goal is not to minimise an evaluation objective on the training dataset, but rather on new, previously unseen samples.

In addition to a training dataset, every machine-learning algorithm needs a hypothesis set, an error function (also called objective function or cost function), and an optimisation procedure. The algorithm searches the hypothesis set to find the hypothesis that best represents knowledge contained in the dataset, which in practice means relying on the optimisation procedure to minimise an estimate of the prediction (i.e., out-of-sample) error. The key here is to strike a balance between the dataset size and the hypothesis-set




Figure 36: Bias-variance tradeoff or how to strike a balance between the hypothesis-set complexity and the dataset size. As the model becomes more and more complex, the training error always decreases. The prediction error, however, decreases only up to a point, and then starts to increase again. The inability of an overly simple hypothesis set to represent the knowledge contained in a dataset creates a bias in predictions, whereas the ability of an overly complex hypothesis set to fit any data perfectly represents noise (i.e., a specific realisation thereof) more than knowledge. In this case, predictions have a large variance.

complexity (Fig. 36). An overly simple hypothesis set contains no single hypothesis that can represent the knowledge that is contained in the data. An overly complex hypothesis set, by contrast, always contains a hypothesis that fits the data perfectly. In doing so, however, the seemingly perfect hypothesis represents not only knowledge, but also the specific realisation of noise, which could be due to genuine stochasticity or measurement errors, or alternatively, deterministic in origin [607].

The performance of an inductively learned model degrades for one of the following three reasons [608]. First, the hypothesis set may not contain a suitable representation of reality. In the case of a classification task, for example, a classifier that is outside of the hypothesis set cannot be learned, although the best hypothesis may still yield a reasonable approximation. Second, the error function may have many local optima over the hypothesis set in which case even representable reality may be hard to learn. Finite data, time,

and memory enable searches through only a tiny subset of all possibilities. Finally, the choice of the optimisation method also determines the scope of search, where methods that try out more hypotheses reduce bias but increase variance and vice versa. Oftentimes it is advantageous to reduce a learning problem to a well-known optimisation problem by transforming the objective function or introducing additional constraints or relaxations. Ref. [609] describes in detail the desirable properties of an optimisation procedure for machine learning. Such properties are good generalisation, scalability to large datasets, good performance in terms of execution times and memory requirements, simple and easy implementation of algorithm, exploitation of problem structure, fast convergence to an approximate solution, robustness and numerical stability for the chosen class of machine learners, and theoretically known convergence and complexity.

Based on feedback available during the learning phase, it is possible to distinguish between three machine-learning contexts: supervised, unsupervised, and reinforcement. In *supervised learning*, the training dataset contains samples that are labelled with an additional ‘ground truth’. A machine learner attempts to learn a target map from data samples to the ground-truth value. If the ground truth is a discrete class from some finite set of classes, then the learner is facing a *classification* task. Sometimes the classification task is such that a probability distribution over the set of classes is more accessible than the direct classifier. If, by contrast, the ground truth is continuous, then the learner is facing a *regression* task.

Among practical obstacles to supervised learning is that obtaining a complete set of labels for all data samples is often difficult and expensive. Learning then uses both a smaller, labelled subset of the full dataset, as well as a larger, unlabelled subset. This type of learning context is referred to as *semi-supervised learning*. The two dominant paradigms in semi-supervised learning are transductive learning and inductive learning [568, 569]. The former does not concern itself with generalisation, but instead attempts to infer the correct labels for the unlabelled subset. This is achieved by assigning labels to unlabelled samples such that ultimately a given optimisation criterion is satisfied across all data, that is, originally labelled and unlabelled subsets taken together. Among the popular techniques are those for learning node representation on networks (i.e., graphs) such as node2vec [503] or DeepWalk [501]. The goal of node-representation learning is to find a low-dimensional space of features with a scalar product that approximates some measure of node similarity in the original network. A commonly used

measure of node similarity is random walks. The learned feature space can be exploited for network community detection, link prediction, and other network-science problems. A more common paradigm, however, is that of inductive learning by which the learner attempts to infer the correct (i.e., generalisable) target map from samples to labels. Irrespective of the employed paradigm, semi-supervised learning is designed to fill the gap between supervised and unsupervised learning, just as the name would suggest.

In *unsupervised learning*, a machine learner is only concerned with extracting insightful patterns from a dataset without relying on any direct feedback. There is no access to supervision signals in the form of discrete or continuous labels. Tasks commonly associated with this machine-learning contexts are partitioning of the dataset into clusters of similar data instances, anomaly detection in the dataset, blind source separation (e.g., picking a single conversation out of many), density estimation of the underlying probability distribution, or learning latent representations of the data by dimensionality reduction techniques. Unsupervised learning is especially useful in exploratory data analysis because of the ability to identify structure in the dataset on its own.

In *reinforcement learning*, a machine learner, often termed an agent, learns how to achieve long-term goals in a complex, uncertain environment. A theoretical underpinning of reinforcement learning is given in the form of Markov decision processes, that is, stochastic decision-making models comprising a set S of agent's states in the environment, a set A of actions available to the agent, and the state-transition probabilities $P(s'|s, a)$ from the state $s \in S$ to a state $s' \in S$ via the action $a \in A$ (Fig. 37). A reward signal r is generated when the agent transitions between states. The objective is to find a policy $\pi(s)$ that prescribes which action should be taken in a given state in order to maximise the cumulative reward obtained over the agent's lifetime. Reinforcement learning comes into play when the environment is so complex or uncertain that the optimal policy is unknowable a priori. The agent learns a model of the environment through execution and simulation, continuously using feedback from past decisions to reinforce good strategies. The learning process is fraught with danger of focusing too much on immediate rewards, thus preventing the discovery of better alternatives, especially those whose rewards are delayed. This is the essence of a trade-off between exploitation and exploration; the agent strives to exploit what is already known in order to accumulate rewards, but at the same time the agent should go exploring to find action selections that pay off more further down the line.

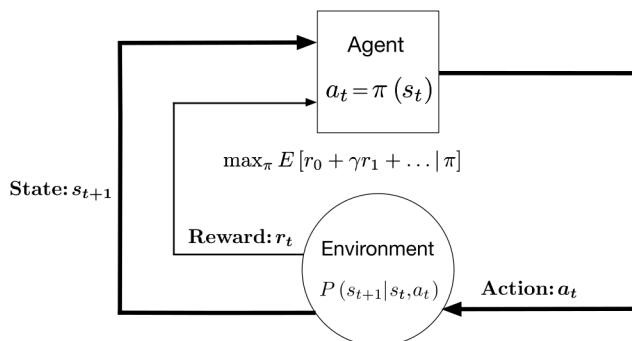


Figure 37: Constituents of a Markov decision process in the reinforcement learning setting. The agent needs to make decisions in discrete rounds $t = 0, 1, \dots, n$. Decisions are made based on a policy $a_t = \pi(s_t)$, which prescribes the action a_t to be taken given the current state s_t . The action triggers a reward signal r_t , while the agent transitions into a new state s_{t+1} according to the transition probability $P(s_{t+1}|s_t, a_t)$. The primary objective is for the agent to find the policy that maximises the total expected reward $E[r_0 + \gamma r_1 + \dots | \pi]$ received over the long run, where γ is a discounting factor. All relevant information about the past is contained in the current state s_t , which encodes aspects of the environment that the agent can sense or influence.

The performance of machine learning approaches discussed heretofore depends heavily on the representation of data. Much effort is therefore put into the feature-engineering process, during which raw data is rendered in a form suitable for modelling. Automating this process is one of the holy grails of machine learning, and will be discussed next.

7.3. Deep learning

Current efforts to automate the feature-engineering process rely on the idea to generate large numbers of candidate features and then select the best feature subsets with respect to a given learning task. This is done while taking into account that features that look irrelevant in isolation may be relevant in combination with other features [608]. Recently, replacing traditional domain expertise and human engineering to hand craft feature extractors has become possible through the development of deep learning [560]. As a subfield of AI and machine learning, deep learning uses multilayer neural networks (hence the term ‘deep’) to exploit the many layers of non-linear information processing for automatic (supervised or unsupervised) feature extraction, and subsequently for pattern analysis and classification (Fig. 38). In deep neural networks, in particular, activation functions are used at the

end of hidden units to introduce non-linear complexities to the model. The most common activation functions are

$$\text{Rectified linear unit} \quad f(x) = \max\{0, x\}, \quad (78a)$$

$$\text{Sigmoid} \quad f(x) = \sigma(x) = \frac{1}{1 + e^{-x}}, \quad (78b)$$

$$\text{Swish} \quad f(x) = x\sigma(\beta x), \quad (78c)$$

$$\text{Hyperbolic tangent} \quad f(x) = \tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}, \quad (78d)$$

$$\text{Softmax} \quad f(\mathbf{x})_i = \text{softmax}(\mathbf{x})_i = \frac{\exp(\mathbf{x}_i)}{\sum_{j=1}^n \exp(\mathbf{x}_j)}, \quad (78e)$$

where softmax is often used for normalising the output layer of a neural network.

Deep learning usually requires large datasets to eliminate the need for manual feature extraction. Because a machine learner is fed with raw data to learn its own representations, deep learning is a form of *representation learning*. Learned representations are contained in the multiple layers of the neural network [610] and encode data by means of a sparse, latent structure with far fewer features than at the beginning. Such elimination of redundant features makes downstream data processing and the final learning task far less intensive. Consequently, most of the recent AI success comes from the utilisation of representation learning with end-to-end trained deep neural-network models in tasks such as image, text, and speech recognition or strategic board and video games. Through enabling the automatic feature engineering, deep learning substantially reduces the reliance on domain-expert knowledge, outperforming in the process traditional methods based on hand-crafted feature engineering, and achieving the performance that equals or supersedes that of humans.

Variants of deep neural networks are designed to improve performance in specific problem domains (Fig. 39). *Convolutional neural networks* thus excel in computer-vision tasks, while *recurrent neural networks* with special gated mechanisms (such as long short-term memory [611] or gated recurrent unit [612]) resolve issues of a vanishing gradient when learning long-term dependencies. Types of encoder-decoder architectures, combined with an attention mechanism [613], are furthermore naturally suited for modelling time series [614] and sequential data, offering state of the art performance in the space of natural-language processing. Finally, the development of




Figure 38: Conceptual illustration of a simple artificial neural network. Generally, artificial neural networks are graph structures comprising multiple layers that perform a number of linear and non-linear transformations on input data. Layers between the first (i.e., input) and the last (i.e., output) layer are called hidden layers. Each layer consists of neurons that receive information from preceding layers across weighted edges. Artificial neural networks propagate information forward to calculate the final output, but also backward to perform weight estimation. *Backpropagation* is the key algorithm that makes weight estimation, and thus the training of deep models, computationally tractable and highly efficient; the algorithm amounts to a shrewd application of the chain rule for derivatives. The gradient of the loss function with respect to each weight is calculated one layer at a time, and then weights are updated in the direction in which the loss function decreases the most. The calculation is iterated until the algorithm converges to accurate outputs on the training dataset.

graph neural networks further accelerated progress in developing general AI architectures for handling unstructured and non-Euclidean data [615].

Despite recent advances in deep learning [616, 617], many obstacles remain to be overcome. The most common drawback is that popular deep-learning techniques need large amounts of data samples in order to generalise and make predictions on unseen inputs, thus being extremely data inefficient. In supervised learning, data inefficiency translates into the need to label, often manually, thousands of data samples; doing so is time-consuming, cumbersome, expensive, and ultimately unreliable. Likewise reinforcement learning demands access to a large number of training trajectories, which in turn must be obtained via human-machine interactions in the real world that are hard to set up. Attempts to resolve the described issues therefore deserve some attention.

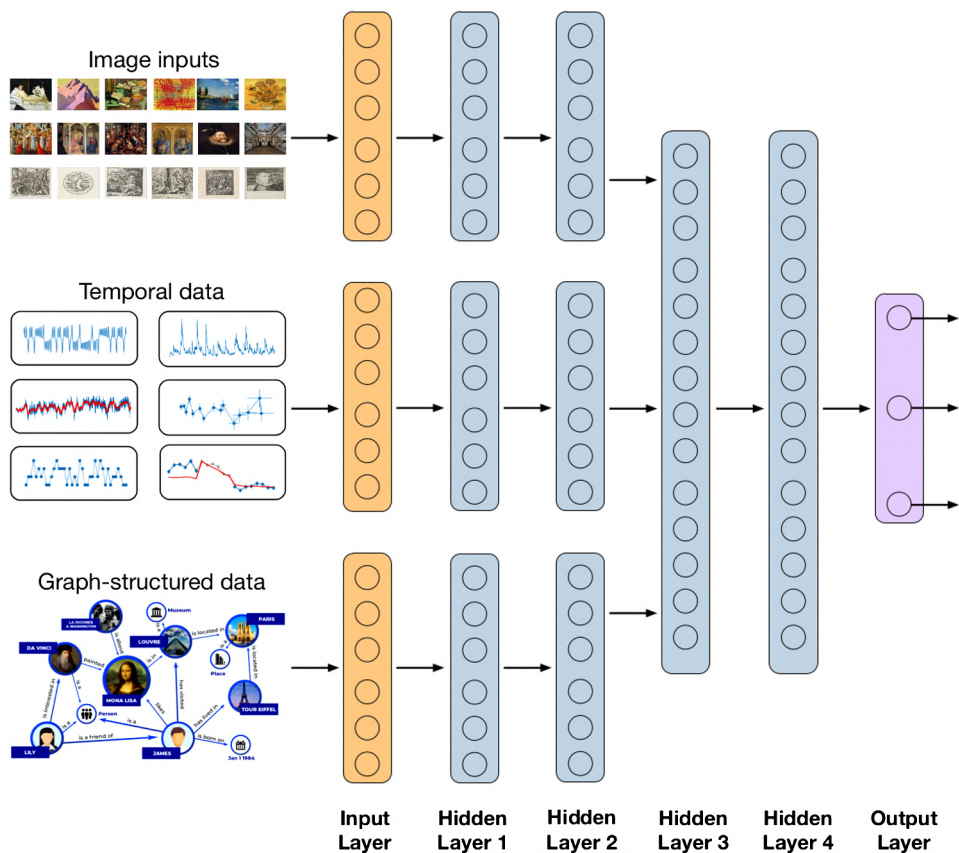


Figure 39: Illustrative example of a large-scale deep neural network. The network accepts as inputs a variety of data types—images, time-series, or graph-structured data—and then in its lower-level hidden layers learns useful representations for each data type.

7.4. Learning to learn

Attempts to improve the data efficiency of deep learning have shown a couple of promising ways forward. *Transfer learning* [618, 619], for instance, relies on the idea that knowledge can be transferred from existing to new models. This approach, inspired by how humans as life-long learning entities use own experiences, exploits structural similarities between learning tasks. To exemplify, an image-recognition model may consist of two parts, a feature-extractor part and a classifier part that is made up of fully connected layers and the output layer. If the model is pretrained for a specific task, the feature-extracting part of the model could still be used for a different task, while fully connected layers of the classifier part are replaced and retrained.

Such retraining requires much less data because only a fraction of weights of the original model must be estimated.

A further step forward towards increased data efficiency, and more fundamentally artificial general intelligence, is *meta learning*. The goal of this subfield of machine learning is to mimic human learning of new concepts [620, 621, 622], which often happens quickly and with only a few examples provided. Meta learning is also known as ‘learning to learn’ thanks to efficiently exploiting previous learning experiences when optimising algorithms to generalise to novel tasks [623]. Such previous experiences include properties of the learning problem, algorithm properties, or patterns already derived from the data, which in turn make it possible to select, alter, or combine elements of learning algorithms to perform well in a previously unseen context.

A common use of meta learning is in the context of supervised few-shot learning [624] which consists of a series of training tasks followed by a series of testing tasks. In a single training task, a dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}$ containing data instances \mathbf{x}_i and their corresponding labels y_i , is divided into a support set \mathcal{S} for learning the task and a query set \mathcal{Q} for determining the classification performance, that is, evaluating the error function. The model parameters are updated based on this performance. Because the support set in each training task contains N different classes with K examples per class, this approach is known as N -way- K -shot classification. Key is that classes differ from one training task to another; to exemplify, let us consider a computer-vision model to distinguish animals of different species. The number of classes may be $N = 3$. One training task for such a model may have the support set with K instances of lions, tunas, and turtles, but the next training task may have the support set with K instances of mice, elephants, and seals. The model tries to use the information in the support sets to classify animals in the query sets. Once the training is complete, testing proceeds on tasks with previously unseen classes, say, cats, dogs, and spiders. The point is that the model learns to discriminate data classes in general (i.e., one animal species from another), rather than a particular subset of classes (e.g., cats from dogs).

Meta-learning approaches can be metric-based, model-based, and optimisation based. Metric-based approaches predict the probability of class y conditional on a data instance \mathbf{x} and the support set \mathcal{S} , which is achieved by means of a weighted sum of labels $y_i \in \mathcal{S}$, where weights are given by a kernel function that measures the distance between the data instance \mathbf{x} and instances $\mathbf{x}_i \in \mathcal{S}$. Well-known examples in this context include Siamese

neural networks [625], matching networks [626], prototypical networks [627], and relation networks [628].

Model-based approaches make no assumptions about the conditional probability of class y ; rather the idea is to design a model for fast learning that updates its parameters over the course of a few training steps. For example, external memory can be used to expedite the neural-network learning process. In the basic setup, a controller neural network receives inputs and generates outputs while reading from and writing to a memory matrix. It is appropriate to think of the controller as a CPU of a computer and of the memory matrix as RAM with a benefit that the whole system can learn to use memory for various tasks instead of sticking to a fixed set of procedures on data. To be usable in a meta-learning context, the described controller-network plus memory-matrix system needs to be trained such that memory encodes information about novel tasks fast and that any stored representation is promptly accessible. Ref. [629] prescribes a training technique that forces memory to hold current inputs for later use. This enables successful classification when a novel instance \mathbf{x} from an already-seen class y is presented at an arbitrary point in time.

Optimisation-based approaches recognise that deep learning models use backpropagation of gradients to learn, and yet the gradient-based optimisation has never been designed to work with a few training samples, nor to converge after a few optimisation steps. To overcome these problems, optimisation itself can be treated as a model to be learned [630]. In the popular model-agnostic meta-learning [631, 632], what is learned is a shared set of model parameter values for initialising optimisation. This shared set leads to quick specialisation on wide variety of tasks, which is achieved by a training procedure that first optimises one shared set of model parameter values to specialise on a batch of tasks, but then uses the results to find an updated shared set that is better at learning with fewer examples.

The described developments have yielded meta-learning methods capable of achieving human and superhuman performance in simple tasks such as one-shot classification. This is just an initial step though. Hopes are that meta-learning approaches will serve an important role in the future discovery of artificial general intelligence.

7.5. AI agents for promoting cooperation

Here, the focus is put on cooperation and AI research aiming to promote cooperation between human and artificial agents in human-machine

networks. Besides human-human (H2H) cooperation already discussed in Section 5, we differentiate between machine-machine (M2M) cooperation, human-machine cooperation facilitated by the former (H2M), and human-machine cooperation facilitated by the latter (M2H).

Artificial learners in human-machine networks are expected to take an active part in society, interacting with both humans and other artificial learners in a complex environment of competition and conflict. What may promote cooperation in such an environment is some form of reciprocity, which underpins the demand for learning algorithms that ensure the emergence of reciprocity in human-machine networks. Evolutionary game theory offers a methodological framework for studying the evolution of cooperation in multi-agent systems in which individual agents must choose between selfish interests and common good. Ref. [633] in particular covers game-theoretical methods in the contexts characteristic of human-machine networks such as crowdsourcing, Internet of Things, and blockchain.

M2M cooperation. In a social-dilemma setting, how can reciprocity, usually observed as a tit-for-tat strategy, emerge in a network of self-interested, reward-maximising reinforcement learners? Ref. [634] shows that naive and commonly defecting reinforcement learners start to cooperate when they incorporate in their own learning process the awareness of their opponent’s learning. Appropriately dubbed learning with opponent-learning awareness or LOLA, the approach leads to the emergence of tit-for-tat and consequent cooperation in the iterated prisoners’ dilemma. LOLA agents exemplify the AI design based on the ‘theory of mind’ [635, 636, 637, 638], that is, the ability to know the opponent’s behaviour and correspondingly alter own behaviour using only human-like, high-level models of other agents rather than the underlying physical mechanisms. Interestingly, agents with a theory of mind about their opponents have a way of dealing with extortionate zero-determinant strategies by being deliberately hurtful until the extortionist opponent becomes fairer [635].

Axelrod’s influential study on the evolution of cooperation [331], involving a round-robin tournament in which strategy entries submitted by game theorists competed in a 200-move iterated prisoner’s dilemma (and which was won by the tit-for-tat strategy), still inspires research today. In fact, there exists a whole Axelrod library [639] of strategies that has been used to organise tournaments similar to Axelrod’s original. Ref. [640] conducted such a tournament with a twist of introducing 5% noise, that is, a chance that an

action is flipped by a random shock. The purpose was to compare the performance and robustness of 176 available strategies for the iterated prisoner’s dilemma. The Axelrod library contains a variety of machine-learning strategies most of which use many rounds of memory, and perform extremely well in tournaments. These strategies encode a variety of other strategies, including the classics such as tit-for-tat, handshake, and grudging. For example, the LookerUp strategy, which does the best in the standard tournament, is a lookup table encoding a set of deterministic responses based on the opponent’s first n_1 moves, the opponent’s last m_1 moves, and the LookerUp agent’s own last m_2 moves. LookerUp is an archetype that can be used to train deterministic memory- n strategies with parameters $n_1 = 0$ and $m_1 = m_2 = n$, which for $n = 1$ cooperate if the last round was mutual cooperation and defect otherwise (known as Grim or Grudger). Interestingly, in this particular tournament, the tit-for-tat strategy could not win any matches.

Pretrained strategies are generally better than human-designed strategies at maximising payoff against a diverse set of opponents. Furthermore, strategies trained using reinforcement learning and evolutionary algorithms with the objective of maximising the payoff difference (rather than own total payoff) resemble zero-determinant strategies, which are generally cooperative and do not defect first, although their performance declines in the presence of noise. Good performance in both standard and noisy tournaments is exhibited by single-layer neural networks albeit their downside is utilising handcrafted features based on the history of play. One of the best-performing strategies in terms of the overall average score is the Desired Belief Strategy [641], which actively analyses the opponent and responds depending on whether the opponent’s action is perceived as noise or a genuine behavioural change. Ultimately, an inescapable conclusion is that reinforcement learning is an effective means to construct strong strategies for various iterated social-dilemma situations [634, 640, 642, 643].

Newer studies go beyond pairwise interactions in the iterated prisoner’s dilemma to examine whether multiple agents cooperate effectively as a team against another team of agents. This requires learning to cooperate with teammates through communication while competing with the opposing team [644, 645, 646]. Ref. [644] proposed an approach called multi-agent deep deterministic policy gradients (MADDPG) that performs well in a number of mixed competitive-cooperative environments. MADDPG is an extension of actor-critic algorithms in reinforcement learning [647, 648]. These algorithms fuse the strengths of actor-only and critic-only methods [649]. The former meth-

ods focus on a parameterised family of policies such that the performance gradient is estimated by simulation, upon which a parameter update is made in a direction of improvement. Among the drawbacks of actor-only methods is that new gradient estimates are independent of past estimates, precluding accumulation and consolidation of previously learned knowledge. Critic-only methods, by contrast, try to estimate the value function (i.e., the total expected reward $E[r_0 + \gamma r_1 + \dots | \pi]$) in order to infer a near-optimal policy from there. A drawback is that there are no guarantees whether the inferred policy will indeed be near-optimal. In MADDPG, the actor is used to select actions, while a central critic evaluates those actions by observing the joint state and actions of all agents. In this sense, MADDPG follows the centralised learning with decentralised execution paradigm [650, 651, 652], which assumes unrestricted communication bandwidth during training, as well as the central controller’s ability to receive and process all agents’ information. To relax these assumptions, Flexible Fully-decentralised Approximate Actor-critic (F2A2) algorithm [653] was proposed as a variant of multi-agent reinforcement learning based on decentralised training with decentralised execution. A strong suit of the F2A2 algorithm is its ability to handle competitive-cooperative partially observable stochastic games [654]. Here, the term ‘partially observable’ refers to situations in which agents only know the probability of making an observation conditional on the current state as opposed to directly determining the current state. ‘Stochastic games’ furthermore refer to situations in which multiple decision makers interact with one another, while the environment changes in response to decisions made.

Hybrid H2M and M2H cooperation. Beside utilising the AI techniques, especially multi-agent reinforcement learning, to learn cooperativeness among machine learners, a growing number of studies investigate hybrid human-machine systems [549, 550, 655, 656]. A key issue in this context is that, in order to interact with humans in social-dilemma situations, machine learners must understand and incorporate moral, trusting, and cooperative human intuitions [657, 658, 643, 659].

Ref. [550] examines how to build machine learners that can cooperate with people and other machines at levels that rival human cooperativeness in two-player repeated stochastic games with perfect information. The study identifies three key properties that algorithms should possess to be successful: (i) *generality* in terms of superior performance in many scenarios rather than a specific one, (ii) *flexibility* to both deter potentially exploitative opponent

behaviours and elicit cooperation in hesitant opponents, and (iii) *learning speed* sufficient to learn effective behaviours after only a few interactions with people. An algorithm displaying these desirable properties is the new simple rule-based expert algorithm termed S++ that uses a version of aspiration learning [660] to select which strategy to follow from a finite set of expert strategies.

An interesting question in the context of hybrid H2M and M2H cooperation is whether humans remain willing to cooperate with machine learners once the true nature of the latter is revealed to the former. In an iterated prisoner's dilemma in which the actions of machine learners were driven by the S++ algorithm [655, 656] shows that cooperativeness goes down when humans assume a non-human opponent. The same happens even in contact with zealous non-human opponents whose behaviour is constant. The results thus point to a transparency-efficiency tradeoff by which being transparent about the true nature of the system is likely to harm efficiency. A possible way around this and similar problems is to combine behavioural- and computer-science expertise to make algorithmic decision-making interpretable by many stakeholders, which in turn would allow people to exercise agency and build trust [661].

Ref. [549] is a particularly nice demonstration that machines can help humans work towards a common goal. In a classic colour coordination game [662] embedded into an artificially-constrained social network, non-human agents proved useful in achieving the collective aim of colouring nodes with one of three colours in such a way that each node's colour differs from the colour of every neighbouring node. Non-human agents used a local optimal colouring strategy with occasional random-colour choices, thus introducing a certain level of noise. Low-noise non-human agents placed centrally in the network improved the resolution of colour conflicts, boosted success rates, and increased the speed with which the problem was solved by nudging humans to occasionally deviate and open up to possibilities. Non-human agents thus facilitated human-human interactions at distant network position in effect helping humans to help each other. Although this illustrative example demonstrates how non-human agents can positively facilitate human cooperation and coordination, there are also examples showing that technology-mediated interactions between humans in online social networks and social-media ecosystems can be used to deceive or manipulate [663, 664, 665, 666]. A task for ongoing and future research is therefore to understand the dynamics of both positive and negative human-machine interactions, and more

importantly offer human-centric solutions the foster the former and avoid the latter.

7.6. Future outlook

AI is expected to enable people to collaborate with machines in an efficient manner for the purpose of solving complex problems. This is further expected to drive the emergence of ever newer and more widespread kinds of human-machine networks. To make the integration of such networks into society as seamless as possible, research heretofore indicates that it will be critical to ensure good mutual communication, trust, clarity, and understanding between humans and machines, that is, the AI technology will have to be human-centric and explainable. Success in integrating human-machine networks into society will then open the doors to a continued massive assimilation of unstructured, semi-structured, or structured data that should be put to good use by addressing a wide spectrum of social-good problems.

Concerns about the black-box nature and opaqueness of deep-learning systems have hampered more widespread AI applications [667]. To address the problem, a strong case is being made for *Human-centric and explainable AI* as a framework towards human-understandable interpretations of algorithmic behaviour. In this way, human operators should be put back into the driver's seat to continually improve the robustness, fairness, accountability, transparency, and explainability [544] of AI technologies.

Furthermore, agreed upon methods to assess the sustained effects of AI on human populations in social, cultural, and political contexts are currently non-existent [667]. Such methods are much needed if AI technologies are to fulfil their promise as an enabler for tackling societal issues or improving human well being. Concrete examples here would be helping to attain sustainable development goals [539] or to alleviate the effects of the Covid-19 pandemic through molecular and clinical breakthroughs [563]. More broadly, but under the condition that the above-stated concerns are resolved, we expect AI to permeate a myriad of research domains and topics, including environmental sustainability (e.g., climate, resource, and biodiversity conservation), social-media abuses (e.g., fake-news, hate-speech, and fraud detection), public safety (e.g., disaster and crime prevention), and more.

8. Criminology

Containing the spreading of crime remains a major challenge across human societies. Empirical data show consistently that crime is recurrent and proliferates, even more so if it is left unchecked. Fig. 40 shows data provided by the United States Federal Bureau of Investigation, indicating that even in strongly monitored and policed states, crime deterrence approaches do not have the desired impact. Indeed, eradicating crime culture is a steep uphill battle, especially in underprivileged social circumstances that do not foster the sense of shared social responsibility. Crime is also problematic in power-driven environments, where greed often overrides the moral compass.

In the realm of physics research, crime is considered as a complex phenomenon, where non-linear feedback loops and self-organisation create conditions that are difficult to foretell, control, and often also difficult to understand [668, 669, 670]. Complexity science in general contends with models in which a large number of relatively simple agents exhibits complex, counter-intuitive, and often unexpected behaviours, and models of crime are in this regard no exception.

Although our understanding of the emergence and diffusion of crime is an ongoing learning experience, recent research shows that methods of statistical physics can significantly contribute to a better understanding of criminal activity. Herein, we review different approaches aimed at modelling and improving our understanding of crime, focusing in particular on the mathematical description of crime hotspots with partial differential equations, on the self-exciting point process and agent-based modelling, adversarial evolutionary games, and the network science behind the formation of gangs and large-scale organised crime. As we hope we will succeed in showing, physics can relevantly inform the design of successful crime prevention strategies, as well as improve the accuracy of expectations about how different policing interventions should impact malicious human activity that deviates from social norms.

8.1. *The broken windows theory*

The 1982 seminal paper by Wilson and Kelling [671] contains many examples and stories that bring the ‘broken windows theory’ to life. For example, how an unattended broken window invites by-passers to behave mischievously or disorderly. Or how a subway graffiti points to an unkempt environment that people can desecrate, signalling also that more egregious

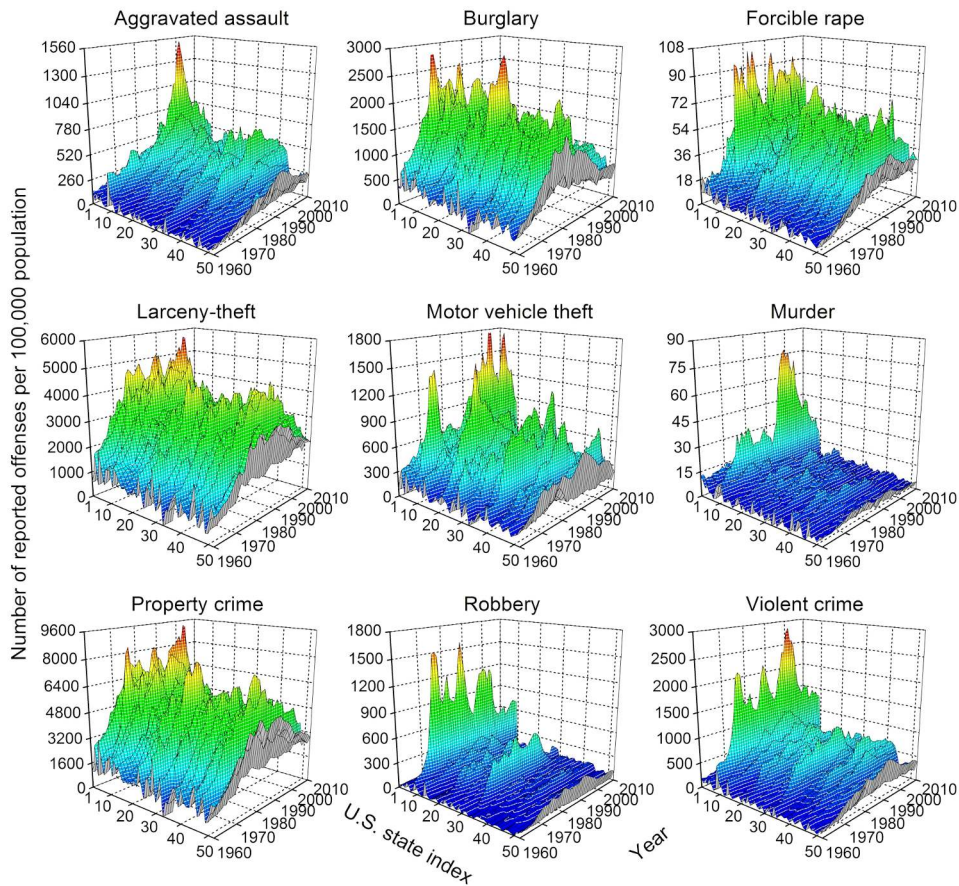


Figure 40: The recurrent nature of crime through lens of data from the Federal Bureau of Investigation. Regardless of type and severity, crime is remarkably recurrent despite our best prevention and punishment efforts. While positive and negative trends are inferrable, crime events (measured as number of offenses per 100,000 population) between 1960 and 2010 fluctuate more or less persistently. More importantly, there is no trend inferable to suggest that crime rates are going down, let alone that crime is vanishing. The U.S. state index is alphabetical, including the District of Columbia being 9th, and the U.S. total being 52nd.

Source: Reprinted figure from Ref. [668].

behaviour might be tolerated. Or how drunks, addicts, prostitutes, and loiterers are more likely to frequent neglected subway stations than orderly and carefully patrolled ones. Thus, on first glance unimportant and petty signals of disorder invite antisocial behaviour and, over time, serious crime—i.e., one broken window soon becomes many.

To physicists, this broken windows theory may be reminiscent of complexity science and self-organised criticality [83], where seemingly small and irrelevant changes at one point in time might have significant and often unexpected and unwanted consequences later on. Moreover, feedback loops, bifurcations, and catastrophes [672], as well as phase transitions [673], are commonly associated with emergent phenomena in complex social systems [674].

Besides the ‘broken windows theory’, there exist other theories of criminal behaviour. According to ‘routine activity theory’ [675], for example, most criminal acts are born out of the convergence of three factors, namely the presence of likely offenders, the presence of suitable targets, and the absence of guardians to protect against the attempted crime. Residential burglary, armed robberies, pick-pocketing, and rape are all examples of such criminal acts.

While intuitively the above three factors are relatively straightforward conditions that obviously favour criminal activity, mathematically they allow us to model the dynamics of criminal offences as deviations from simple random walks. This is due to built-in heterogeneities in target selection that may drive criminal activity towards preferred locations and away from less desired ones. The degree of target attractiveness may change in time and depend on mundane factors such as the day of the week or weather conditions, or on the more sophisticated interplay between landscape, criminal activity, and law enforcement responses. Crime dynamics may also include learning mechanisms or feedback loops. These elements ultimately lead to the emergence of non-trivial patterns such as spatially localised crime hotspots [676] and repeat and near-repeat victimisation [677, 678, 679, 680], wherein the odds of a second victimisation of the original target or a target in its vicinity are greatly enhanced.

The complexity of crime dynamics that stems from the above-described fundamental considerations has as a consequence the fact that the mitigation and displacement of crime is a highly non-trivial task—a task which, based for example on data shown in Fig. 40, we often fail at [681, 682, 683, 684, 685].

In this light, it is important to note that straightforward gain-loss principles that underlie rational choice theories are likely too simple and naive. Anticipating that stronger punishment would just lead to less crime is simply not aligned with the reality, not in empirical data, and not from mathematical models that at least to some degree attempt to capture the complexity of crime [686, 687].

8.2. Crime hotspots

Presented empirical observations of spatio-temporal clusters of crime in urban areas (Fig. 41) motivated the development of a statistical model of criminal behaviour [676]. The model was developed to study residential burglary, where target sites are stationary, which is not that case in crimes where offenders and targets are mobile, as in assault or pick-pocketing. The model builds on the assumption that burglars are opportunistic, and that they thus victimise areas that are sufficiently close to where they live, and where possibly they have committed crimes before [688]. Another important assumption is that the distances that criminals are willing to travel to engage in criminal acts are best described by monotonically decreasing functions [689]. The movement of offenders is usually described as a biased random walk, whereby the bias is twofold as follows. In the first place, a given home may be intrinsically more attractive to a burglar due to its perceived wealth, the ease of access, or the predictable routine of its residents. Secondly, there may be learned elements that bias the burglar towards a specific location, for example to a previously victimised home where a successful break-in was once already possible.

To quantify the bias towards any given location and to determine the subsequent rate of burglary, the hotspot crime model includes a dynamically changing attractiveness field [676]. Moreover, the tendency for repeat victimisation is included in the model by temporarily increasing the attractiveness field in response to past crimes [690, 691]. Since residential burglary entails non-moving crime targets, and for simplicity, it is convenient to start with a discrete model on a square lattice with periodic boundary conditions. Each lattice site $s = (i, j)$ is a real estate with attractiveness $A_s(t)$ and the number of criminals $n_s(t)$. The higher the value of $A_s(t)$, the higher the bias towards site s and the more likely it will be the subject of crime. Moreover, once site s has been victimised, its attractiveness further increases. The following decomposition is introduced

$$A_s(t) = A_s^0 + B_s(t), \quad (79)$$

where A_s^0 is the static, though possibly spatially varying, component of the attractiveness field, and $B_s(t)$ represents the dynamic component associated with repeat and near-repeat victimisation. More precisely, $B_s(t+1) = B_s(t)(1 - \omega) + E_s(t)$, where ω sets a time scale over which repeat victimisations are most likely to occur, while $E_s(t)$ is the number of events that




Figure 41: Dynamic changes in residential burglary hotspots in Long Beach California, as observed for two consecutive three-month periods, starting in June 2011. The emergence of different burglary patterns is related to how offenders move within their environments and how they respond to the successes and failures of their illicit activities. Returns to previously victimised locations or locations in their vicinities are common and in agreement with the ‘routine activity theory’ [675]. Source: Reprinted figure from Ref. [676].

occurred at site s between t and $t + 1$. To take into account the broken windows theory [671], we let $B_s(t)$ spread locally from each site s towards its nearest neighbours s' according to

$$B_s(t + 1) = \left[(1 - \eta)B_s(t) + \frac{\eta}{z} \sum_{s'} B_{s'}(t) \right] (1 - \omega) + E_s(t) \quad (80)$$

where the sum runs over the nearest neighbour sites associated to site s , z is the coordination number of the lattice, and η is a parameter between zero and one that determines the significance of neighbourhood effects. Higher values of η lead to a greater degree of spreading of the attractiveness generated by a given burglary event, and vice-versa for lower values. For simplicity, we can further assume that the spacing between sites ℓ and the discrete time unit δt over which criminal actions occur are both equal to one, and that every time

a site s is subject to crime its dynamic attractiveness $B_s(t)$ increases by one. Interaction networks other than the square lattice, which better describe the city grid or social networks can be easily accommodated as well.

Criminal activity is included in the model by allowing individuals to perform one of two actions at every time step. A criminal may either burglarise the site they currently occupy, or move to a neighbouring one. Burglaries are modelled as random events occurring with probability $p_s(t) = 1 - \exp[-A_s(t)]$. Whenever site s is subject to crime, the corresponding criminal is removed from the lattice, representing the tendency of actual burglars to flee the location of their crime. To balance this removal, new criminal agents are generated at a rate Γ uniformly at random on the lattice. If site s is not burglarised, the criminal will move to one of its neighbouring sites with probability $1 - p_s(t) = \exp[-A_s(t)]$. The movement is thus modelled as a biased random walk so that site s' is visited with probability

$$q_{s \rightarrow s'}(t) = \frac{A_{s'}(t)}{\sum_{s'} A_{s'}(t)}, \quad (81)$$

where the sum runs over all neighbouring sites of s . The position of the criminals and the biasing attractiveness field in Eqs. (79) and (80) create non-linear feedback loops which give rise to complex patterns of aggregation that are reminiscent of actual crime hotspots, similar to those depicted in Fig. 41. The model actually displays four different regimes of $A_s(t)$, as shown in Fig. 3 of Ref. [676], all of which apply to different realities of residential burglary.

A continuum version of the above-described discrete model has also been introduced [692, 693], the bifurcation analysis of which can also outline suggestions for crime hotspot suppression and policing. According to [676], the continuum version of the dynamics of the attractiveness field takes the form

$$\frac{\partial B}{\partial t} = \frac{\eta D}{z} \nabla^2 B - \omega B + \epsilon D \rho A, \quad (82)$$

where $D = \ell^2/\delta t$, $\epsilon = \delta t$, and $\rho(s, t) = n_s(t)/\ell^2$. The continuum equation for criminal number density, denoted as ρ , is given by

$$\frac{\partial \rho}{\partial t} = \frac{D}{z} \vec{\nabla} \left[\vec{\nabla} \rho - \frac{2\rho}{A} \vec{\nabla} A \right] - \rho A + \gamma, \quad (83)$$

where offenders exit the system at a rate ρA , and are reintroduced at a constant rate per unit area $\gamma = \Gamma/\ell^2$. Eqs. (82) and (83) are coupled partial




Figure 42: Failure and success of crime hotspot suppression. In the upper row, crime hotspots emerge as a supercritical bifurcation. When subjected to suppression, they simply displace but never vanish completely. New hotspots always emerge in positions adjacent to the original ones. In the lower row, crime hotspots emerge via a subcritical bifurcation. When subjected to suppression, the hotspot gradually vanishes without giving rise to new hotspots in nearby locations. The colour maps encode the time evolution of the attractiveness field B . We refer to Ref. [692] for further details. Source: Reprinted figure from Ref. [692].

differential equations that describe the spatio-temporal evolution of the attractiveness B and the offender population ρ , and they belong to the general class of reaction-diffusion equations that frequently exhibit spatial pattern formation [694].

In order to study the effects of police intervention, the crime rate ρA in Eq. (83) is set to zero at given hotspot locations and for a given time frame [692]. Calculations then show that only subcritical crime hotspots may be permanently eradicated by means of a suitable suppression mechanism, while supercritical hotspots are only displaced but never fully removed from the population (Fig. 42).

The mathematical models describing the nucleation and diffusion of crime

hotspots can be upgraded to include spatial disorder, as well as approaches to dynamically adapt suppression measures to evolving crime patterns, or to choose from different deployment strategies and more rigorous analysis [695, 696, 697, 698, 699]. Along similar lines, related research includes the consideration of dynamical systems that take into account the competition between citizens, criminals, and guards [700], the effects of socio-economic classes and changes in police efficiency and resources allocated to them [701], the impact of imprisonment and recidivism [702], and the possibility of self-defence of communities against crime [703].

8.3. *Crime as a self-exciting point process*

Certain types of crime exhibit similar space and time clustering as earthquake activity. Examples include burglary, gang violence, and property crime. Just like clustering patterns observed by seismologists indicate that the occurrence of an earthquake is likely to induce a series of aftershocks near the location of the initial event, so are these types of crime likely to reoccur near initially victimised spots, thus leading to crime swarms and clusters, and lending themselves to the application of seismology methods to model criminal activity. The self-exciting point processes is one such method [704].

A space-time point process is defined by a collection of points with location (x, y) at time t , where a certain event took place. This event can be an earthquake, a lightning strike, or a criminal act. The process is then described by a conditional rate $\lambda(x, y, t)$, which gives the occurrence rate at location (x, y) in dependence on the history $H(t)$ of the point process up to time t [705]. There is typically an initial or parent crime, akin to a parent earthquake, which generates several follow-up or offspring crimes, akin to aftershocks. The follow-up crimes are described by a triggering function $g(x, y, t)$, which depends on previous criminal activity, but with an amplitude that decreases with increasing spatio-temporal distance from it. In modelling crime, a multiplicative factor $\nu(t)$ for the background activity is also needed, which takes into account fluctuations due to weather, seasonality, or day time. Decades of research in seismology have lead to well-defined forms for above functions. In crime, on the other hand, non-parametric methods and calibrations using data are necessary for their estimation. For details we refer to the seminal work by Mohler et al. [704].

The self-exciting point process has been applied and tested on urban crime using residential burglary data from the Los Angeles Police Department [704]. Traditionally, crime hotspot maps were generated by means of

a pre-assigned fixed kernel, using previous crime occurrences as input [706]. However, the point process methodology has been found to yield better results, and this also for types of crime where near-repeat effects do not play such a prominent role, like robberies and car theft. The main source of this superiority has been attributed to a better balance between exogenous and endogenous contributions to crime rates and to the method relying on direct inference from data, rather than on an imposition of hotspot maps using a pre-assigned fixed kernel.

Self-exciting point processes have also been used to analyse temporal patterns of civilian death reports in Iraq [707]. Similarly to urban crimes, the rate of violent events has been partitioned into the sum of a Poisson background rate and a self-exciting component in which previous bombings or other episodes of violence generate a sequence of offspring events according to a Poisson distribution. The study showed that point processes are well suited for modelling the temporal dynamics of violence in Iraq.

The geographic profiling of criminal offenders can also be made using self-exciting point processes in order to estimate the probability density for the home base of a criminal who has committed a given set of spatially distributed crimes. Target selection from a hypothetical home base is informed by geographic inhomogeneities such as housing types, parks, freeways or other physical barriers, as well as directional bias and preferred distances to crime [708]. These techniques have also been used to model intra-gang violence that results from retaliation after an initial attack [699].

Future research along this line could be aimed at further refining point process models towards crime type and local geography. Doing so would facilitate the application of this promising methodology.

8.4. Social dilemmas of crime

The prisoner's dilemma game is amongst the most frequently employed theoretical frameworks to study pairwise social dilemmas [709]. In the prisoner's dilemma game, two players should simultaneously decide whether to cooperate or defect, and then based on their choices receive payoffs accordingly. A social dilemma arises because mutual cooperation yields the highest collective payoff, but the payoff for a defector is higher if the opponent decides to cooperate. Mutual defection is therefore the only rational outcome if we assume that both players act in self-interest so as to maximise their individual payoffs. In the long run this leads to the proliferation of defection

and ultimately to the ‘tragedy of the commons’ [710, 34], where common resources are lost to societies due to overexploitation and lack of shared social responsibility.

While criminal behaviour does not map directly to the prisoner’s dilemma game, the framework of evolutionary games, and evolutionary social dilemmas in particular [234, 291], lends itself very well to modelling crime. In this context, social order can be considered as the public good that is threatened by criminal activity, with competition arising between criminals and those trying to prevent crime. However, committing crimes is not necessarily equivalent to defection, because unlike defectors, criminals may actively seek to harm others. Likewise, fighting crime is often more than just cooperating, in particular because it may involve risk that goes beyond contributing some fraction of one’s wealth into a common pool. Thus, a more deliberate formulation of competing strategies may elevate, and is in fact needed for, the accuracy of the modelling approach.

With these considerations in mind, an adversarial evolutionary game with four competing strategies, as shown in Fig. 43, has been proposed in Ref. [711]. The game entails informants (I) and villains (V) as those who commit crimes, as well as paladins (P) and apathetics (A) as those who do not. Informants and paladins actively contribute to crime abatement by collaborating with authorities whenever asked to do so. All players may witness crimes or be the victims of crime, in agreement with victimisation surveys [712]. Thus, paladins are model citizens that do not commit crimes and collaborate with authorities. At the other end of the spectrum we have the villains, who commit crimes and do not report them. Somewhere in between we have informants who report on other offenders while still committing crimes, and apathetics who neither commit crimes nor report crimes of others. The lack of active cooperation in apathetics may be due to inherent apathy, fear of retaliation, or ostracism from the community at large. Apathetics are similar to second-order free riders in the context of the public goods game with punishment [713, 714], in that they cooperate at first order by contributing to the public goods as in not committing crimes, but defect at second order by not punishing offenders.

At each round of the game a criminal is selected randomly from the $V + I$ pool together with a potential victim from the $N - 1$ remainder of the population. The two selected players begin the game with a unitary payoff. After a crime occurs, the criminal player increases their payoff by δ , while the victim loses δ . If the victim is either an apathetic or a villain, the crime is




Figure 43: Crime as a four-strategy evolutionary game, comprising informants, paladins, villains, and apathetics. The four strategies are defined by their propensities to commit crimes and serve as witnesses in criminal investigations. Arrows between strategies indicate the number of possible game pairings and outcomes in which the update step leads to a strategy change. For example, there are two ways by means of which a villain can be converted into a paladin. Circular arrows within each strategy quadrant indicate updates such that player strategies remain unchanged.

Source: Reprinted figure from Ref. [711].

not reported to the authorities and therefore successful: the victim's payoff is decreased to $1 - \delta$ and the victimiser's is increased to $1 + \delta$. If, on the other hand, the victim is a paladin or an informant, the crime is reported to the authorities and an investigation begins. For this, a subset M of the $N - 2$ remaining players is drawn, and the victimiser is convicted with probability $w = (m_P + m_I)/M$, where m_P and m_I are the number of paladins and informants within M . In case of a conviction, the victim is refunded δ , and the payoff of the criminal becomes $1 - \theta$, where θ determines the punishment fine. With probability $1 - w$ the crime is left unpunished, in which case the criminal retains $1 + \delta$, while the victim's payoff is further decreased to $1 - \delta - \epsilon$, where ϵ is due to retaliation of the accused who, having escaped punishment, feels empowered in their revenge. Other interpretations of ϵ may be damages

to personal image or credibility, or loss of faith in the system after making an accusation that is unsubstantiated by the community. Notably, in the latter case, the choice of reporting one's victimisation to authorities may be even more detrimental to the witness than the original criminal act ($\epsilon > \delta$), which is common in societies that are heavily marred by war, by mafia, or drug cartels, where very few people will serve as witnesses to crimes.

Parameter values of δ , θ , and ϵ are always used such that all payoffs remain positive. At the end of each round of the game, the player with the smaller payoff changes its strategy according to proportional imitation [715]. In particular, if the victimiser is emulated, the loser simply adopts the victimiser's strategy and ends the update as either a villain or an informant. If the victim is emulated, the loser mimics the victim's propensity to serve as a witness but adopts a noncriminal strategy regardless of the victim's. In this case, the update results with the loser becoming either a paladin or an apathetic (see Fig. 43 for details).

Simulations of the four-strategy evolutionary game described above reveal that informants are key to the emergence of utopia—a crime-free society. Indeed, a crime-dominated society can become crime-free by imposing an optimal number of informants I_0 at the onset of the game. The dynamics depends on the chosen parameter values. A utopia may be elusive in extremely adversarial societies in which initially we have high numbers of villains and apathetics. However, by deriving a deterministic version of the above described game [711], it is possible to show that if there are at least some informants initially present in the population ($I_0 > 0$), the final state is always utopia regardless of δ , θ , and ϵ (Fig. 44).

While beneficial, the presence of informants may come at a cost, either in training an undercover informant, or in convincing a criminal to collaborate with authorities, or in tolerating the criminal acts that informants will keep committing. One may thus consider an optimal control problem [716] to investigate the active recruitment of informants from the general population in terms of associated costs and benefits. Higher recruitment levels may be the most beneficial in abating crime, but they may also be prohibitively expensive. The optimal control problem was expressed via three control functions subject to a system of delay differential equations. The research showed that optimal recruitment strategies change drastically as parameters and resource constraints vary, and moreover, that more information about individual player strategies leads only to marginally lower costs [716].

The important role of informants within the reviewed adversarial evolu-




Figure 44: The emergence of utopia in a society with informants. All trajectories with $I_0 > 0$ evolve towards a crime-free state. The ternary diagram shows unstable fixed points in light red, unstable fixed lines in thick light red, stable fixed lines in thick dark blue, and trajectories beginning (or ending) along various eigenvectors as thick green arrows. The dystopian fixed point d and the saddle point s are unstable to increases in I , so that the only attracting final states for $I_0 > 0$ are those utopias with $P > P_c$. These results were obtained with $\delta = 0.3$, $\theta = 0.6$, and $\epsilon = 0.2$. Source: Reprinted figure from Ref. [711].

tionary game [711] has also been studied by means of human experiments in [717]. The goal was to test whether, and if yes to what degree, informants are actually critical for crime abatement as predicted by theory. Remarkably good agreements between simulations and laboratory experiments have been observed for different parameterisations of the game, thus lending full support to the approach.

In addition to social dilemmas, the evolution of crime can also be studied by means of the inspection game [718]. Rational choice theories predict that increasing fines should diminish crime [686]. However, a three strategy inspection game in which, in addition to criminals (C) and punishing inspectors (P), ordinary individuals (O) are present leads to significantly different and counterintuitive outcomes [719, 720]. The O players neither commit crimes nor do they participate in inspection activities. They represent the masses that catalyse rewards for criminals and costs for inspectors. Ordinary individuals receive no bonus payoffs upon encountering inspectors or their peers. Only when paired with criminals do they suffer the consequences of crime in

form of a negative payoff $-g \leq 0$. Criminals, on the other hand, gain the reward $g \geq 0$ for committing a crime. When paired with inspectors, criminals receive a payoff $g - f$, where $f \geq 0$ is a punishment fine. When two criminals are paired none of them receive any benefits. Inspectors, on the other hand, always have the cost of inspection, $c \geq 0$, but when confronted with a criminal, an inspector receives the reward $r \geq 0$ for a successful apprehension. This game was studied via Monte Carlo simulations on a square lattice with periodic boundary conditions where each lattice site is occupied either by a criminal, a punishing inspector, or an ordinary citizen. The game evolves by first randomly selecting player s to play the inspection game with their four nearest neighbours, yielding the payoff P_s . One of the nearest neighbours of player s , s' , is then chosen randomly to play the game with their four nearest neighbours, leading to $P_{s'}$. Finally, player s' imitates the strategy of player s with probability

$$q = \frac{1}{1 + \exp\left(\frac{P_{s'} - P_s}{K}\right)}, \quad (84)$$

where K determines the level of uncertainty in the strategy adoption process. The chosen form in Eq. (84) corresponds to the empirically supported multinomial logit model [721], which for two decision alternatives becomes the Fermi function [722, 302]. A finite value of K accounts for the fact that better performing players are readily imitated, although it is possible to adopt a strategy by player who is performing worse, for example due to imperfect information or errors in decision making.

Monte Carlo simulations reveal that the collective behaviour of the three-strategy spatial inspection game is indeed complex and counterintuitive, with both continuous and discontinuous transitions between different phases. Here a phase is either a single-strategy or a multi-strategy stable state that is unavoidable by any other combination of strategies or a single strategy. Usually, evolutionary games with more than two competing strategies require the stability of subsystem solutions be performed for the accurate and correct determination of phase transitions [723]. A subsystem solution can be formed by any subset of all possible strategies. The winner between two subsystem solutions can be determined by the average moving direction of the invasion front that separates them, yet it is crucial that the competing subsystem solutions are characterised by a proper composition and spatio-temporal structure before the competition starts. In this way the three-strategy inspection game also yields a cyclic dominance phase [259], which emerges spontaneously due

to pattern formation and is robust against initial condition variations, and in which all three competing strategies coexist.

Taken together, these results indicate that crime should be viewed not only as the result of offending actions committed by certain individuals, but also as the result of social interactions between people who adjust their behaviour in response to societal cues and imitative interactions. The emergence of crime thus should not be ascribed merely to the ‘criminal nature’ of particular individuals, but also to the social context, the systems of reward and punishment, the level of engagement of the community, as well as to the interactions between individuals. This more comprehensive view of crime may have relevant implications for policies and law enforcement.

8.5. Criminal networks

The goal of this section is to review how methods of physics, and in particular of network science [439, 724, 422, 425, 725, 282, 426, 273], can contribute to better understanding organised crime [726], such as drug cartels, the formation of gangs, or political corruption networks [727].

Criminal structures like the Italian Mafia [728], street gangs, or drug cartels [729] often emerge when fear and despair become so ingrained within a society that the social norm is simply to accept crime. In such a case, witnesses and even victims of crime often choose not to cooperate with law enforcement in the prosecution of criminals. Instead, people sometimes try to fit in, although acquiescence and acceptance are usually slippery slopes towards later forms of active engagement. Ultimately this thus leads to the growth of a criminal network.

Criminological research has identified a number of factors that may promote the regional development of crime, including unemployment [730, 731], economic deprivation [732], untoward youth culture [733], failing social institutions [734], issues with political legitimacy [732], as well as lenient local law enforcement strategies [735, 736], to name the most prominent examples. Policies aimed at reducing recruitment into organised crime have also been incorporated in agent-based models with a multiplex-network structure to capture the effects of household, kinship, school, work, friends, and co-offending social relations [737]. Recent work on declining criminal behaviour in the U.S. suggests that trends in the levels of crime may be best understood as arising from a complex interplay of a rich myriad of said factors [738, 739], while most recent empirical data indicate that social networks of criminals

have a particularly strong impact on the occurrence of crime—the more the criminals are connected into networks, the higher the crime rate [740, 741].

The assumption that there is a network structure behind organised crime invites the idea that removing the leader, or the most important hubs of the network [742], will disrupt the organisation sufficiently to hinder or at least heavily disrupt criminal activity. Law enforcement agencies thus often attempt to identify and arrest the ‘ring leader’ of an identified criminal organisation. But even if successful, such operations rarely have the desired effect. A recent study analysing cannabis production and distribution networks in the Netherlands shows that this strategy may in fact be fundamentally flawed [743]. All attempts towards network disruption analysed in the study proved to be at best unsuccessful (Fig. 45). At worst, they had the opposite effect in that they have increased the efficiency of the network. The latter was achieved by means of nifty reorganisations, such that the attack ultimately made these networks stronger and more resilient to future such attempts. By combining computational modelling and social network analysis with unique criminal network intelligence data from the Dutch Police, Duijn et al. [743] have concluded that criminal network interventions are likely to be effective only if applied at the very early stages of network growth, before the network gets a chance to organise, or to reorganise to maximum resilience.

Gang rivalries have been studied by means of agent-based simulations in conjunction with data from the Hollenbeck policing division of the Los Angeles Police Department [744]. The details of the model were as follows. Each agent is part of an evolving rivalry network that includes past interactions between gang members. Individuals perform random walks where the jump length is drawn from a truncated Lévy distribution and where bias in the direction of rivals is included. Gang home bases, historical turfs, and geographic details that may limit movement such as freeways, rivers, and parks were all taken into account in the simulated biased Lévy walk network. Typical gang behaviour, as inferred from the criminology literature, has also been considered. Using metrics from graph theory, it was possible to show that simulated biased Lévy walk network modelling is in fact the most accurate in replicating actual gang networks. In Fig. 46, we reproduce a picture from [744], showing simulated results and an actual map of violent crimes in Hollenbeck, which are indeed in very good agreement. This approach can also be used to infer unknown rivalry interactions, in particular because the simulated biased Lévy walk network converges to stable long-term configura-




Figure 45: A criminal network in the Netherlands involved in cannabis cultivation. Nodes represent the actors that are needed for successful production and distribution of cannabis. The network is highly resilient to targeted disruption strategies. Even worse, research shows that perturbations may lead to reorganisation towards an even more robust and resilient network. Node sizes represent the number of actors fulfilling the associated role, and link thickness corresponds to the total number of links between actor groups. Source: Reprinted figure from Ref. [743].

tions. The authors of Ref. [744] have also noted that the method is portable and can be applied to other geographical locations, offering insight into gang rivalry distributions even in the absence of data. The method may also be extended to test sociological concepts related to gang interactions such as territoriality and allegiances within gangs.

Police department field interview cards were later used to study the behavioural patterns of roughly 748 suspected gang members who were stopped and questioned in Hollenbeck [745]. The goal was to identify any social communities among street gang members by creating a fully-connected ad hoc

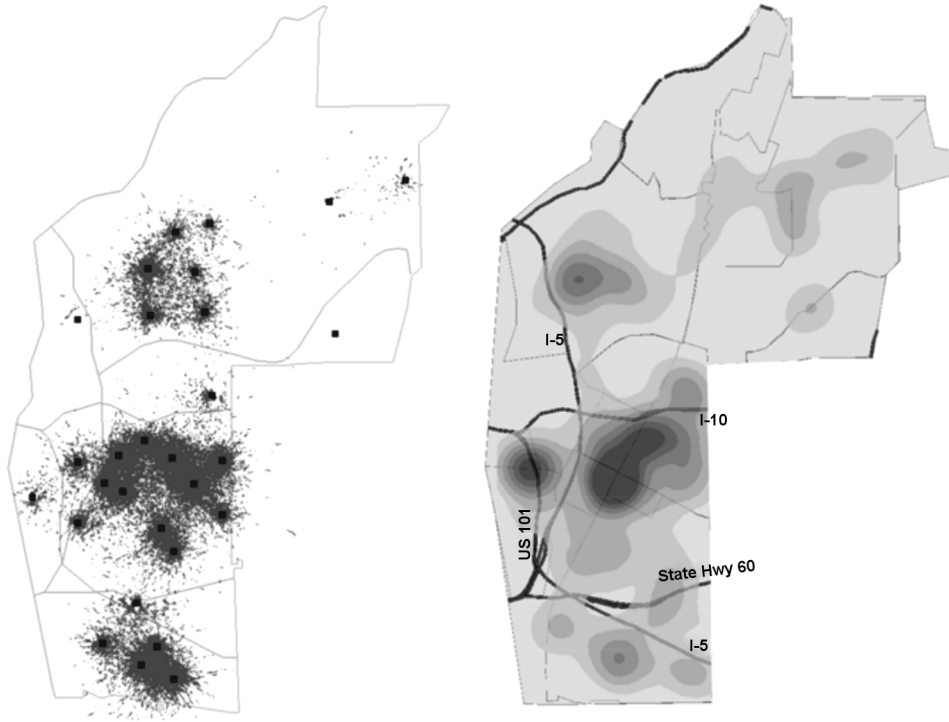


Figure 46: Reconstructing a gang network from data with the biased Lévy walk network method. Interactions between agents simulated using the biased Lévy walk network method are shown left, while the actual density map of gang-related violent crimes in Hollenbeck between 1998 and 2000 is shown right. Thick lines represent major freeways crossing the city.

Source: Reprinted figure from Ref. [744].

network where individuals represent nodes and links encode geographical and social data. Individuals stopped together were assumed to share a friendly or social link and the distance $d_{i,j}$ between stop locations of individuals was recorded. This information was used to determine the affinity matrix $W_{i,j}$ associated with the network. Its entries are composed of a term that decays as a function of $d_{i,j}$, representing geographical information, and of an adjacency matrix whose entries are zero or one depending on whether individuals were stopped together or not. The latter represents social information. By using spectral clustering methods, distinct groups were then identified and interpreted as distinct social communities among Hollenbeck gang members. These clustered communities were finally matched with actual gang affilia-

tions recorded from the police field interview cards.

To evaluate the quality of identified clusters, the authors of Ref. [745] used a purity measure, defined as the number of correctly identified gang members in each cluster divided by the total number of gang members. Results showed that using geographical information alone leads to clustering purity of about 56 % with respect to the true affiliations of the 748 individuals taken in consideration. Adding social data may improve purity levels, especially if this data is used in conjunction with other information, such as friendship or rivalry networks. These results may be used as a practical tool for law enforcement in providing useful starting points when trying to identify possible culprits of a gang attack.

An interesting physics-inspired approach to modelling gang aggregation and territory formation by means of an Ising-like model has also been proposed in Ref. [746]. In particular, otherwise indistinguishable agents were allowed to aggregate within two distinct gangs and to lay graffiti on the sites they occupy. Interactions among individuals were indirect and occurred only via the graffiti markings present on-site and on nearest-neighbour sites. Within this model, gang clustering and territory formation may arise under specific parameter choices, and a phase transition may occur between well-mixed and well separated, clustered configurations. In the mean-field version of the model, parameter regimes were identified where the transition is first or second order. In all cases, however, these clustering transitions were driven by gang-to-graffiti couplings since direct gang-to-gang interactions were not included in the model. The role of graffiti and vandalism has been reviewed also by Thompson et al. [747], who analysed the urban rail industry, where graffiti markings have significant impact on expenditure, timely operation of services, and on passenger perception of safety.

Methods of network science are also well suited to study a more subtle form of crime, namely political corruption. Indeed, corrupt behaviour in politics limits economic growth [748, 749, 750, 751, 752], embezzles public funds [753], and promotes socio-economic inequality in modern democracies [750, 754]. The World Bank estimates that the annual cost of corruption exceeds 5 % of the global Gross Domestic Product, which amounts to \$2.6 trillion USD, with \$1 trillion USD being paid in bribes around the world. In another estimation by the non-governmental organisation Transparency International, corrupt officials receive as much as \$40 billion USD bribes per year in developing countries, and nearly two out of five business executives have to pay bribes when dealing with public institutions. Despite the diffi-

culties in trying to estimate the cost of global corruption, there is consensus that massive financial resources are lost every year to this cause, leading to devastating consequences for companies, countries, and the society as a whole.

The shortage of studies aimed at understanding the finer details of corruption processes is in considerable part due to the difficulties in finding reliable and representative data about people who are involved [755]. On the one hand, this is certainly also because those who are involved do their best to remain undetected, but also because information that does leak into the public is often spread over different media outlets offering conflicting points of view. In short, lack of information and misinformation [756] both act to prevent in-depth research.

To overcome these challenges, Refs. [727, 757, 758] have employed datasets that allow in-depth insights into corruption scandals in Brazil and Mexico. The Brazilian dataset [727] in particular provides details of political-corruption activities of 404 people who were from 1987 to 2014 involved in 65 important and well-documented scandals. Notably, Brazil has been ranked 79th in the Corruption Perceptions Index, which surveyed 176 countries in its 2016 edition, which places it behind African countries such as Suriname (64th) and Ghana (70th), and way behind its neighbouring countries such as Uruguay (21th) and Chile (24th). Methods of time series analysis and network science have been applied to reveal the dynamical organisation of political corruption networks in Brazil, which in turn reveals fascinating details about individual involvement in particular scandals, and it allows the prediction of future corruption partners with useful accuracy [727]. Research showed that the number of people involved in corruption cases is exponentially distributed, and that the time series of the yearly number of people involved in corruption has a correlation function that oscillates with a four-year period. This indicates a strong relationship with the changes in political power due to the four-year election cycle. By linking together people that were involved in the same political scandal in a given year, it was also possible to create a network representation of people that took part in corruption scandals (Fig. 47). The network has an exponential degree distributions with plateaus that follow abrupt variations in years associated with important changes in the political powers governing Brazil. By maximising the modularity of the latest stage of the corruption network, we can observe statistically significant modular structures that do not coincide with corruption cases but usually merge more than one case.

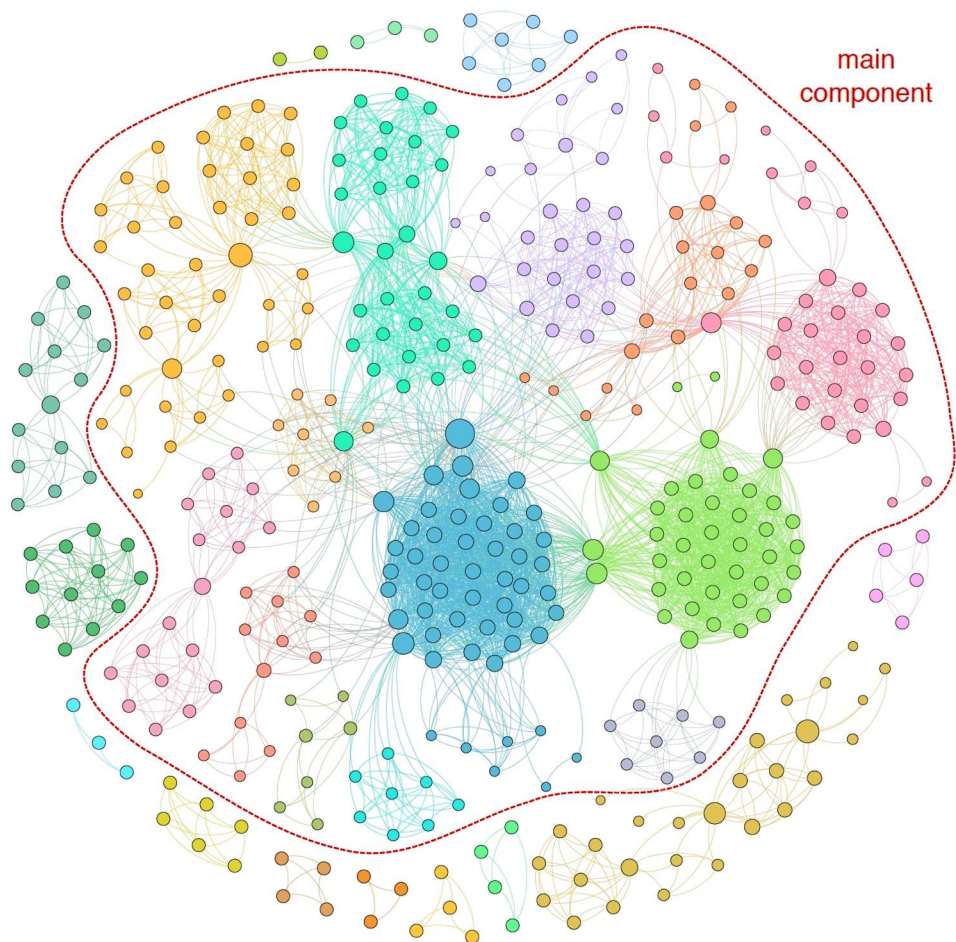


Figure 47: Network representation of people involved in corruption scandals in Brazil from 1987 to 2014 (from Ref. [727]). Each vertex represents a person and the edges among them occur when two individuals appear at least once in the same corruption scandal. Node sizes are proportional to their degrees and the colour code refers to the modular structure of the network, as obtained with the network-cartography approach [759]. There are 27 significant modules, and 14 of them are within the giant component indicated by the red dashed loop.

Source: Reprinted figure from Ref. [727].

Based on this research, it is also possible to apply different algorithms for predicting missing links in corruption networks. By using a snapshot of the network in a given year, Ribeiro et al. [727] have tested the ability of these algorithms to predict missing links that appear in future iterations of the

corruption network. Obtained results show that some of these algorithms have a significant predictive power in that they can correctly predict missing links between individuals in the corruption network, which could be used effectively in prosecution and mitigation of future corruption scandals.

Lastly, we mention promising efforts to detect criminal organisations [760] and to predict crime [761] based on demographics and mobile data. It is known that the usage of communication media such as mobile phones and online social networks leaves digital traces, and research shows that this data can be used successfully for detecting and characterising criminal organisations. We hope that this section shows that, with the help of network science and community detection [423, 424], law enforcement agencies could better understand hierarchies within criminal organisations, more reliably identify members who play central roles in them, as well as obtain valuable information on connections among different sub-groups and their respective responsibilities in the illicit undertakings.

8.6. Rehabilitation and recidivism

The final, and perhaps even the most important, stage in treating crime is the rehabilitation of past offenders. Only if past offenders acknowledge and understand their wrongdoing, and only if after paying their dues they can be integrated successfully back into society, can we consider the problem solved. Otherwise, we are patching the problem of crime with temporary solutions that in the long run do not actually lead to better societal outcomes. Sadly, the available data indicate that it is in these later stages of crime abatement where societies fail the most, and in particular fail to give offenders a second chance at a new start in life, which often pushes them, or at least provides a nudge, into recidivism.

The dilemma that commonly shows in such cases is often referred to as the ‘stick versus carrot’ dilemma. In other words, should rehabilitation programs focus on punishing wrongdoing (stick), or should they focus on generously rewarding steps of progress along the way (carrot)? There is ample research in evolutionary game theory that addresses this dilemma in the context of cooperation in the public goods game [34]. Notably, there is no simple conclusion or resolution of the dilemma, due to the fact that the outcome depends substantially on the context and other circumstances that are taken into account in the model. Theoretically at least, punishment seems to be more promising simply because it can stop once the target behaviour is achieved. Rewarding, on the other hand, often creates a self-enforcing

loop in that the more progress is achieved, the higher the rewards that are expected to uphold the good trend. However, research on punishment also emphasises the very negative consequences of antisocial punishment and with it related concerns to use sanctions as a means to promote collaborative efforts and to raise social welfare [762, 763]. Evidence suggesting that rewards may be as effective as punishment and lead to higher total earnings without potential damage to reputation [764] or fear from retaliation [318] has also been mounting. Moreover, Rand and Nowak [765] provide firm evidence that antisocial punishment renders the concept of sanctioning ineffective, and argue further that healthy levels of cooperation are likelier to be achieved through less destructive means.

Regardless of whether the burden is placed on punishment [766, 767, 768, 714] or reward [769, 770, 771, 772, 773, 774] or both [775], the problem with both actions is that they are costly. Cooperators who abstain from either punishing or rewarding therefore become second-order free riders, and they can seriously challenge the success of sanctioning as well as rewarding. In the context of rehabilitating criminals, the question is how much punishment for the crime and how much reward for eschewing wrongdoing in the future is in order for optimal results, as well as whether these efforts should be placed on individuals or institutions [776, 777], all the while also assuming of course that the resources are limited [778, 779].

To improve our understanding of these important considerations, Berenji et al. [780] have introduced an evolutionary game to study the effects of carrot and stick intervention programs on criminal recidivism. Their model assumes that each player may commit crimes and may be arrested after a criminal offence. In the case of a conviction, a criminal is punished and later given resources for rehabilitation in order to prevent recidivism. After their release into society, players may choose to continue committing crimes or to become paladins (P). The later option is an optimal outcome, indicating they have been permanently reformed. Players are given r chances to become paladins. If after the r -th arrest and rehabilitation an individual relapses into crime it is marked as an unreformable (U). States P and U are sinks, meaning they mark the end of the evolutionary process for each particular individual. The P/U ratio is therefore a natural order parameter, such that societies with a lot of crime are characterised by $P/U \rightarrow 0$ while crime-free societies are characterised by $P/U \rightarrow \infty$. The main parameters of the game are the resources allocated for rehabilitation h , the duration of the rehabilitation τ , and the severity of punishment θ .

Simulations of this model have been performed which include the constraint $h\tau + \theta = C$, where C is the total amount of available resources, and where $h\tau$ is the part of these resources that are spent on rehabilitation—the carrots—while θ is the remainder, spent on punishment—the sticks. Because C is finite, increasing one effort decreases the other, hence the ‘stick versus carrot’ dilemma. As C increases, the ratio P/U will increase as well (Fig. 48). This means that with more general resources available, the conversion to paladins becomes more efficient. For a given value of C , the most successful strategy in reducing crime, warranting the highest P/U ratio, is to optimally allocate resources so that after being punished, criminals experience impactful intervention programs, especially during the first stages of their return to society. Indeed, the upper right panel of Fig. 48 reveals that for the case of $N = 400$ players, the optimal parameter values are $h = 0.3$, $\tau = 1.5$ and $\theta = 0.35$, which indicates that the available resources C need to be balanced so that there is enough stick (a sufficiently high θ) and enough carrots (a sufficiently high h) for a long enough time (a sufficiently high τ). Within this model, excessively harsh or lenient punishment is less effective than when the two are well-balanced. In the first case, there are not enough resources for rehabilitation left, in the second, punishment was not strong enough to discourage criminals from committing further crimes upon release to society.

The findings reviewed in this section have important sociological implications, and they provide useful guidance on how to minimise recidivism while maximising social reintegration of criminal offenders. At the same time, we note that research dedicated specifically to rehabilitation and recidivism at the interface of physics and criminology is rather sparse, so that this is certainly an avenue worth exploring more prominently in the future, especially given its importance in assuring long-term success of prior crime prevention strategies.

8.7. *Rosy outlooks for less crime*

The physics of crime is a developing and vibrant field, with ample opportunities for novel discoveries and improvements of existing models and theory. The model of crime hotspots, for example, could be upgraded to account for the distribution of real estate that better reflects the layout of an actual city. It would then be interesting to learn whether and how the introduced heterogeneity in the interaction network affects the emergence and diffusion of hotspots. If the crime is no longer residential burglary but




Figure 48: Minimising recidivism requires carefully balanced rehabilitation programs, where both punishment and reward play a crucial role. Either neglecting punishment in favour of generous rehabilitation or vice versa will ultimately fail in successfully reintegrating offenders into society. Depicted is the ratio between paladins and unreformables P/U in dependence on the amount of resources for rehabilitation h , as obtained for different values of the duration of intervention τ (see top of individual graphs). In all cases the severity of punishment θ is adjusted so that $h\tau + \theta = C$ (see legend in the top left graph), taking into account the fact that available resources are finite. The upper right graph reveals that the optimal parameter values are $h = 0.3$, $\tau = 1.5$, and $\theta = 0.35$, which indicates that the most successful strategy is to allocate the limited resources so that after being punished, criminals experience impactful intervention programs, especially during the first stages of their return to society.

Source: Reprinted figure from Ref. [780] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

crime that involves moving targets, further extensions towards social networks whose structure varies over time also become viable. If crime is treated as an evolutionary game the possibilities for upgrades range from increased strategic complexity to the integration of more realistic, possibly co-evolving, interaction networks that describe human interactions. In the realm of adversarial evolutionary games, it would be interesting to study the impact of different strategy adoption rules, in particular because imitation-based rules

are frequently contested with best-response dynamics in the realm of human behaviour. In addition to the outlined extensions and upgrades of existing models, it is also possible to envisage new classes of models, especially such that would build more on self-organisation and growth from first principles to eventually arrive at model societies with varying levels of crime. Here the hierarchical growth of criminal networks involving persuasion to join an organisation and fidelity to either committing or not committing crimes could be fertile starting grounds.

As we hope this section shows, the physics of crime can provide useful insights into the emergence of criminal behaviour, as well as indicate effective policies for crime mitigation. We also hope the reviewed results may be useful to police and other security agencies in developing better and more cost-effective crime mitigation schemes while optimising the use of their limited resources. Indeed, the physics of crime has far-reaching implications, and we emphasise that the time is ripe for these insights to be used in synergy with traditional crime-related research to yield more effective crime mitigation policies. Many examples of ineffective policies clearly highlight that an insufficient understanding of the complex dynamical interactions underlying criminal activity may cause adverse effects from well-intended deterrence strategies. A new way of thinking, maybe even a new kind of science for deterring crime is thus needed—in particular one that takes into account not just the obvious and similarly linear relations between various factors, but one that also looks at the interdependence and interactions of each individual and its social environment. One then finds that this gives rise to strongly counterintuitive results that can only be understood as the outcome of emergent, collective dynamics. This is why physics can make important and substantial contributions to the better understanding and containment of crime.

The aim here is to highlight valuable theoretical resources that can help us bridge the gaps between data and models of criminal activity. Employing these resources should certainly contribute to rosier outlooks for human societies with less crime.

9. Migrations

Incoming and outgoing migrations, respectively called immigration and emigration, pose substantial challenges to society. Unsuccessful integration of immigrants, for instance, often leads to cultural and socio-economic seg-

regation that, if unchecked, may trigger unrest and ethnic clashes. The explosiveness of such situations reveals that interethnic tolerance is subject to cascades and tipping points as the harbingers of radical political transformations. Once a tipping point is crossed, tolerance typically evaporates fast on its own, but the transformation can further be catalysed by shocks in the form of economic crises or pandemics. Fortunately, concepts and methods studied in statistical physics and its relatives, complexity and network sciences, can help develop our quantitative understanding of large-scale social dynamics. Examples in this context include, but are not limited to, tipping points and phase transitions [673, 781, 782], cascade failures [266, 783], resilience or robustness [742, 784], and recoveries or repairs [785, 786].

When migrations take place, numbers matter. The European Union (EU), for example, handled in an orderly manner about 300,000 asylum seekers yearly up until 2014, but as that number quadrupled in two short years (Fig. 49), a deep crisis emerged, prompting prominent political figures to prognosticate an end to the EU as a political project [787]. Some members of the Schengen Zone responded to the crisis by invoking the 'exceptional circumstances' provision of the Article 26 of the Borders Code to unilaterally reinstate internal border controls, while others chose to erect barbed-wire fences along borders with their non-Schengen neighbours. Without a clear solution in sight, the debate on the subject polarised around two 'ideological blackmails'; one side argued that the EU's borders must stay fully opened to refugees, whereas the other side argued that the borders must be swiftly and completely closed [788]. The crisis furthermore served as a platform for the rise of right-wing political populism throughout Europe [535]. That a large increase in migrant inflows can provoke such a knee-jerk reaction in terms of strengthening border controls and embracing populist policies suggests conditional—albeit heterogeneous—tolerance levels towards diversity of peoples, values, lifestyles, etc. Indeed, the general population seems to heed Karl Popper's maxim that unconditional tolerance must lead to the demise of tolerance [789].

In the language of thermodynamics, and as evidenced by the existence of international borders, human society is far from an equilibrium state. Maintaining a steady non-equilibrium state requires a costly investment of energy, which is precisely the cost that the European countries attempted to avoid by forming the EU and the Schengen Area. The recent migrant crisis, however, is a stark reminder that freeing trade and the movement of people demands utmost care to harmonise relations not only between political elites




Figure 49: Main corridors and destination countries during the European migrant crisis. The crisis started in 2014, peaked in 2015, and was declared over in 2019. Germany received by far the largest number of asylum requests in 2015, followed by Hungary, Sweden, Austria, Italy, and France (vertical bars). Source: Reprinted figure from Ref. [790].

who orchestrate agreements, but also on the microscopic scale of individual interactions. It is in this latter context that the economist Durlauf argued that statistical mechanics can inform research in social sciences [791]. The basic idea of statistical mechanics—that every atom is influenced by other atoms even beyond just the immediate neighbours—is similar to the ideas of social science that an individual’s decisions depend upon the decisions of others. This, in turn, has led to an intriguing possibility that a common mathematical formalism underlies natural and social phenomena [792, 793, 794].

When people migrate, be it for personal safety (refugees) or in search of better life (economic migrants), they are forced to establish new contacts and friendships, as well as acclimatise to a new language and culture as a part of an integration process [795]. Individual socio-economic interactions that take place during this process often pose social dilemmas that involve balancing selfish interests and common good [225, 796]. Evolutionary game theory offers a formal framework to resolve social dilemmas, bringing into focus five mechanisms of social viscosity that help ‘lubricate’ interactions between individuals. These mechanisms include three types of reciprocities,

direct, indirect, and network, and two types of selection, kin and group [234]. Intriguingly, evolutionary game theory has maintained a close tie to statistical physics ever since the discovery of game-driven spatial chaos in structured populations [244, 797]. Later, this has led to an even stronger tie to network science, especially via the emergence of network reciprocity [242] and subsequent discoveries that put social networks at the forefront of resolving key social dilemmas [250, 301].

We have, heretofore, identified statistical physics, complexity and network sciences, and evolutionary game theory as some of the fields that could help model large-scale social dynamics due to migrations. Incidentally, the concept of phase transitions has exerted tremendous influence on all these fields [798, 799, 800, 801], begging the question as to what causes such widespread fascination with this concept. It has become increasingly evident that, aside from physical systems such as water and ferromagnetic materials, many dynamical complex systems also possess critical thresholds—called tipping or crossover points—near which the system undergoes a swift transition from one state to another [802]. Among the famous examples of this are ecosystems [803], but similar arguments have been made for stock-market crashes and political revolutions [804, 805].

In the wake of the EU’s migrant crisis, and in view of momentary flirting with border controls and populist policies during the crisis, it is worthwhile examining if the EU members approached a tipping point sometime in 2015 or 2016. Although a devastating phase transition that would put a stop to the EU political project has been averted for now, some signs of a system transitioning from one state to another have remained in public records. For instance, by the second half of 2015, an estimated 58% of the EU citizens harboured concerns about immigration, more than a double compared to a year before (24%) and more than a triple compared to two years before (16%) [806]. A similar rise of a single overwhelming concern had previously been seen in the second half of 2011, when economic situation preoccupied an estimated 59% of the EU citizens, and thus paved the way for a landmark surrender of the European Central Bank to ultra-low interest-rate policies. More important, however, is the fact that both in 2011 and 2015, a normally multidimensional space of a population’s concerns shrank to the point when one dimension dominated all. This shared-concern phenomenon has many parallels with long-range order in physical systems by which a system’s remote constituents exhibit highly correlated behaviour [807]. The ordered state is often established via symmetry breaking upon a phase transition

from a disordered state, with a famous example of this being the spontaneous magnetisation of ferromagnetic materials below the Curie temperature [808]. Aside from physical systems, long-range order in terms of the cross-correlation between remote constituents has been observed in biological systems, concretely, the healthy operation of cellulo-social collectives [809]. Natural and social sciences have thereafter made progress by generalising the described ideas to the notion of self-organisation. This notion has proven relevant to human endeavours on all scales ranging from economics [810] to robotics [811] to traffic flows [124] and many others [812].

Returning to the problem of migrations, the EU crisis had in the meantime abated, especially when the focus shifted to the ongoing Covid-19 pandemic [566]. The underlying causes of migrations, however, have not been resolved (e.g., Middle Eastern turmoils and African armed conflicts [813]). Adding to this state of affairs the predicted consequences of climate change, future migrant waves are to be expected [814]. The question is then is the current world order ready?

9.1. Tolerance

A way to define the notion of tolerance, or toleration, is mutual acceptance of conflicting worldviews without resorting to suppressive violence. The worldview of others is in this notion seen in a negative light, yet there are vindicating circumstances that outweigh the negatives [815, 816]. Immigration squarely fits into this ‘tolerance dichotomy’ because immigrants are commonly seen as a threat to job security and a source of increased competition in the job market, especially during economic downturns, yet immigrants also serve as a much needed workforce in labour-deficient sectors, especially during the periods of economic growth [817]. Aside from the economic dimension, the space of factors that shape attitudes towards immigrants has multiple other dimensions. These can roughly be classified as individual or collective [818]. The former include educational attainment, cultural conflicts, political involvement, interpersonal trust, and public safety, whereas the latter include immigrant abundance and foreign investments.

That there exists a plethora of factors affecting attitudes towards immigrants goes a long way in confirming that tolerance is conditional rather than unconditional. In Germany, for example, a sudden increase in the immigrant inflow at the onset of the migrant crisis tightly correlated with the increase in the popularity of an anti-immigrant right-wing populist party (Fig. 50). An even starker example is that in nine out of 15 polled Eastern European

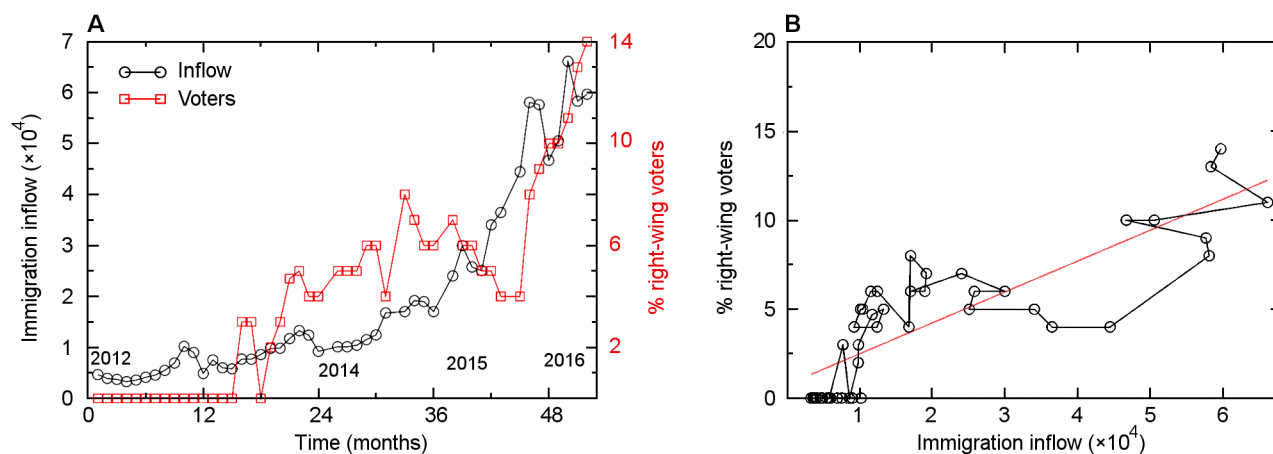


Figure 50: Right-wing populist parties gain popularity when the immigrant inflow increases. **A**, In Germany, the increasing inflow of immigrants rather clearly coincided with the increasing support for a right-wing populist party. **B**, Significant regression emerges when the German case is presented as a scatter plot between the immigrant inflow and the percentage of right-wing populist voters.

Source: Reprinted figure from Ref. [535] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

countries in 2016, more than half the population expressed views that their country should refuse any Syrian refugees [819]. Interestingly, none of these nine countries have been the main refugee destination, and yet their appetite for immigration was minute compared to that of their Western European neighbours, suggesting a strong cultural divide in tolerance. The described conditional nature of tolerance in the real world is in sharp contrast with primary legislation, specifically the European Convention on Human Rights, which envisions unconditional tolerance by stating that any refugee or displaced person has the right to EU protection if an asylum is claimed within the EU. It is this latter notion of tolerance that Karl Popper challenged [789], as do simple models of the evolution of human cooperation. In the model of Nowak and Sigmund [820], cooperation evolves via indirect reciprocity as practised by conditional cooperators, whereas unconditional cooperators inadvertently undermine cooperation because they help others regardless of how those others behave, thus giving defectors an edge to invade and prevail. Similarly in the tag-based model of cooperation by Riolo, Cohen, and Axelrod [821], an overly tolerant population is vulnerable to mutants who rarely give help to others.

The conditional nature of tolerance, and especially its manifestations such as anti-immigrant sentiments and right-wing populism [822, 806, 823, 535, 824], indicate that we need a better quantitative understanding of the interplay between the processes of immigration and integration. To this end, Ref. [825] studies the balance between immigration and integration rates in relation to the tolerance of the local population for newcomers. Combining the elements of evolutionary game theory and network science, an artificial society is envisioned to form a social network that grows not just intrinsically, but also by attracting newcomers from the outside. Newcomers are attracted to a benefit differential gained by cooperating with the locals. All cooperative interactions, represented by social-network links, yield a per-capita fitness Φ_1 for locals and Φ_2 for newcomers such that the growth of the two sub-populations is given by $N_1(t+1) = \Phi_1(t)N_1(t)$ and $N_2(t+1) = \Phi_2(t)N_2(t)$. Denoting the fraction of newcomers in the population with $f_n = N_2/N$, where $N = N_1 + N_2$ is the population size, we have that $f_n(t+1) = R(t)f_n(t)$. Here, $R = \Phi_2/\bar{\Phi}$ is the newcomer-to-population fitness ratio and $\bar{\Phi} = (1 - f_n)\Phi_1 + f_n\Phi_2$ is the average population fitness. As long as the benefit differential is positive, it holds that $R > 1$, causing the fraction of newcomers in the social network to grow. The benefit differential is a government policy that cannot be changed by individuals which is why the following individual-scale processes take place. First, as locals get increasingly surrounded by newcomers, the tolerance of the former for the latter gradually saturates. Second, tolerance saturation causes a local to either radicalise and stop cooperating with newcomers altogether, or to remain cooperative but tacitly support radicals. Third, depending on their surroundings, newcomers can pick up the cultural patterns of locals, thus becoming integrated and no longer exerting pressure on the tolerance of locals. The social dynamics arising from the described setup, as will be explained next, distinguishes between a successful and an unsuccessful immigration policy.

Outcomes of a given immigration policy can be predicted analytically using a mean-field approximation or estimated numerically using computer simulations. Examples of predictable outcomes are the probability $\Pr(X)$ that a randomly chosen local is radicalised or the probability $\Pr(Y)$ that the number of radicals and their tacit supporters exceeds a critical threshold (N_1^c) necessary to initiate oppressive action against newcomers. In a democracy, this might be electing a right-wing populist party to power. The mean-field

analysis shows that

$$\Pr(X) = \frac{1}{k_{\max} - k_{\min} + 1} \sum_{l=k_{\min}}^{k_{\max}} \sum_{k=l}^{\langle k \rangle} \binom{\langle k \rangle}{k} f_n^k (1 - f_n)^{\langle k \rangle - k}, \quad (85)$$

where k_{\min} (respectively, k_{\max}) is the number of newcomer neighbours at which the least (resp., the most) tolerant local radicalises, whereas $\langle k \rangle$ is the average number of neighbours (i.e., the average node degree). The mean-field analysis further shows that

$$\Pr(Y) = e^{-\lambda} \sum_{l \geq N_1^c} \frac{\lambda^l}{l!} \approx 1 - F_{\text{norm}} \left(\frac{N_1^c}{N_1}; \frac{\lambda}{N_1}; \frac{\sqrt{\lambda}}{N_1} \right), \quad (86)$$

where $\lambda = N_1 f_n (1 - g_r)$ and g_r is the fraction of radicalised locals. The function F_{norm} is the cumulative distribution function of a normal random variable with the mean $\mu = \lambda/N_1$ and the standard deviation $\sigma = \sqrt{\lambda}/N_1$. The mean-field and numerical results are in good agreement (Fig. 51).

Ultimately, the phase space of immigration policies can be explored using numerical simulations. By doing so, three types of outcomes reveal themselves:

- *Mutualism* is a set of equilibrium states reached by a smooth reduction of the fitness ratio R to a level at which locals maintain a sustainable majority (continuous blue curve in Fig. 52A).
- *Newcomer dominance* is a set of equilibrium states, also reached by a smooth reduction of the fitness ratio R , but to a level at which newcomers form a majority (continuous blue curve in Fig. 52A).
- *Antagonism* is a set of non-equilibrium, absorbing states due to a complete breakdown of cooperation between the two sub-populations (discontinuous red curve in Fig. 52A).

If society is sufficiently tolerant and integrates newcomers at a reasonable rate, it is likely for local and newcomer populations to interact in a mutually beneficial manner (Fig. 52B). If, however, integration is too slow, a peaceful transition to the newcomer majority may take place (Fig. 52B). Finally, peace may succumb to turmoil and violence if slow integration is matched with low tolerance (Fig. 52B). These outcomes show that a successful immigration policy must be carefully thought out, taking into account (i.e., measuring and monitoring) the tolerance of locals and the integrability of newcomers.

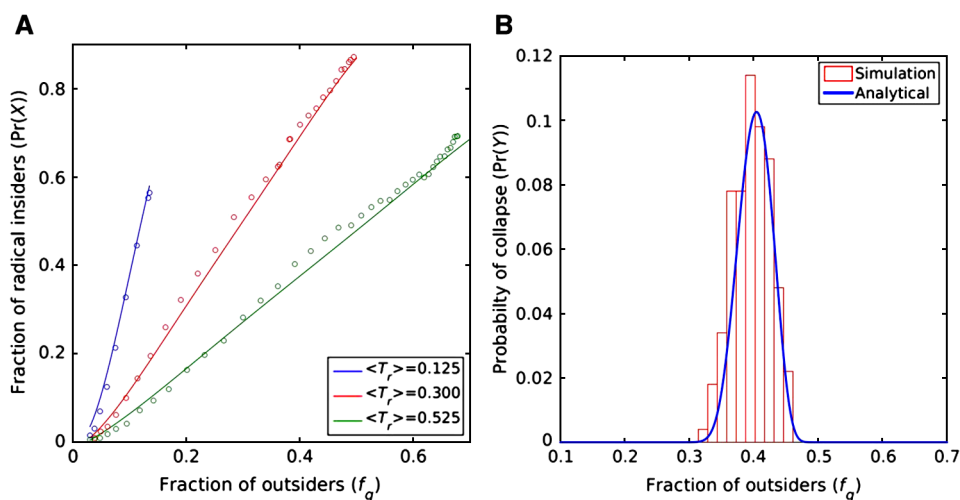


Figure 51: Outcomes of immigration policies. **A**, Probability of a randomly selected local being radicalised is a monotonically increasing function of the fraction of newcomers. The tolerance of the local population for newcomers controls how strongly the former react to the latter, as well as the maximum abundance of the latter. **B**, Probability of reaching the critical number of radicals and their tacit supporters to attain a majority need not always increase with the fraction of newcomers. If the tolerance of the local population is sufficient, tolerant locals together with newcomers comprise the majority. Source: Reprinted figure from Ref. [825].

9.2. Integration and culture

Defining an exact criterion that would mark the end of the integration process is difficult. To partly circumvent the problem and quantitatively analyse the situation, economists often look at economic integration that compares the earnings of immigrants to those of natives. Early cross-sectional analysis showed that at the time of arrival in the United States immigrants earn 17% less than natives, but the difference disappears in 15 years, whereas after 30 years immigrants even earn 11% more than natives [826]. Borjas, however, argued that the cross-sectional picture fails to account for a gradual decline in skills of immigrant cohorts after 1965, concluding that almost all immigrants since the second half of the sixties experience “the same sluggish relative earnings growth” and that earnings parity between immigrants and natives is “extremely unlikely” [827].

Cultural integration is particularly difficult to express in quantitative terms. Some anecdotal evidence paints a picture that integration is facilitated by ‘cultural similarity’ between native and immigrant populations.

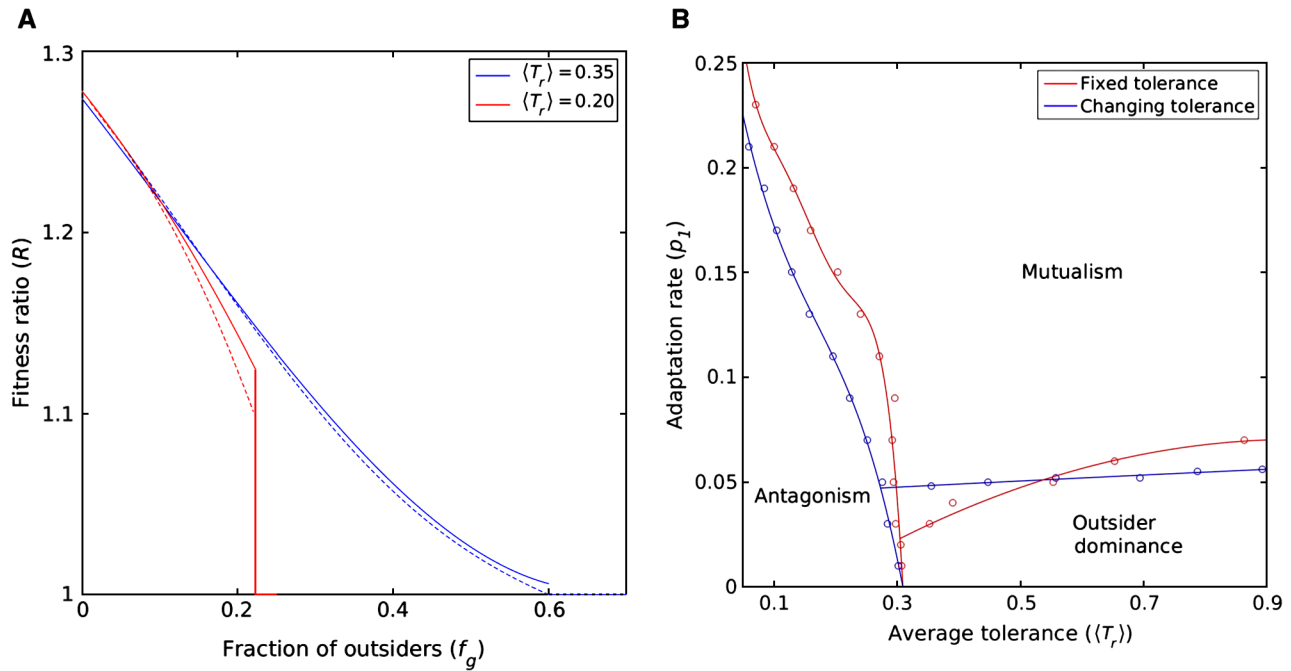


Figure 52: Immigration-policy outcomes. **A**, Newcomer-to-population fitness ratio R decreases because abundant newcomers saturate the tolerance of locals who then stop being cooperative. If locals are relatively tolerant, the decrease in R is continuous. By contrast, if locals are relatively intolerant, the decrease in R may be discontinuous, signifying a sudden collapse of cooperation between newcomers and locals. Solid curves are simulations, whereas dotted curves are a mean-field approximation. **B**, Success of immigration policies depends on the tolerance of locals and the integration rate of newcomers. Tolerant locals and integrable newcomers are likely to interact in a mutually beneficial manner. Tolerant locals, however, may lose the majority to non-integrable newcomers. Finally, intolerant locals and non-integrable newcomers make for an explosive combination. Source: Reprinted figure from Ref. [825].

Largely successful immigration policies in this context are considered to be those of Australia and Canada, whereas within the EU borders, there are Luxembourg, Portugal, and Spain. In Luxembourg, for example, 45% of the total population are foreign nationals, who over 90% originate from other EU countries, meaning that the cultural similarity between native and immigrant populations is relatively large. If Luxembourg is too small and too rich to be taken as a representative example, there is the case of Portugal, where somewhat over 3% of the population is foreign and dominated by the Brazilians (who are ‘cultural neighbours’ of the Portuguese). The 3% num-

ber, however, is deceptively small because Portugal has naturalised many of its immigrants, illustrating the importance of successful integration of the migrant population. Furthermore, in Spain, slightly under 10 % of the population comprises foreign nationals, again dominated by cultural neighbours from Latin American countries. Interestingly, there is also a sizeable proportion of culturally more distant minorities, and Spain has even been targeted by terrorist attacks, yet the right-wing populist movement has never gained a foothold. The Spanish case thus illustrates the importance of a balanced approach which secures peaceful and prosperous co-existence of a highly diverse population.

It is instructive to contrast the above-mentioned successful cases with the situation in France. The French population comprises about 8.5 % of foreign nationals, but also an additional 10.5 % direct descendants of immigrants. About half of these are of culturally distant origin, such as Arab-Berber, Sub-Saharan, and Turkish. Difficulties in absorbing such a large and culturally distant migrant population prompted a series of terrorist attacks across France to which the native majority responded by increasingly swinging towards the political right. Ultimately, the French case is a sort of antithesis to the Spanish one, and illustrates just how hard it may be to strike the much needed balance when dealing with an inflow of migrants.

The cultural-similarity hypothesis finds some empirical support in the data on economic integration as well. Being of non-EU origins and living in a non-mixed household (i.e., having a non-native spouse) had a significant negative impact on immigrant earnings in a range of EU countries [828]. This negative effect was smaller for immigrants with non-EU origins who lived in a mixed household (i.e., with a native spouse) or for EU-born immigrants who lived in a non-mixed household. Finally, the effect was insignificant for EU-born immigrants who lived in a mixed household. From a theoretical perspective, the concept of cultural similarity was introduced by Axelrod in the context of his seminal model of social influence and cultural change [829].

Axelrod's model of social influence is based on three principles:

1. *Agent-based modelling* means that mechanisms are first specified at the individual scale, and then the consequences of such mechanisms are examined at the population scale to discover the collective or emergent properties of the system.
2. *The lack of a central authority* means that cultural change occurs in a bottom-up manner without coordination from a global overseer.

3. *Adaptive instead of rational actors* means that local circumstances dictate how influencing or being influenced takes place (cf. Ref. [830]). There is no cost-benefit analysis nor forward-looking strategic evaluation.

Culture in Axelrod's model is a multidimensional trait set. Each cultural dimension (say, formal wear) is accompanied by its own traits (say, morning dress, dress suit, ceremonial dress, uniform, religious clothing, national costumes, or frock coats). In an abstract form, a trait is represented by an ordinal number, implying that culture itself is a list of trait numbers. If two actors have the same culture, then all their corresponding trait numbers are equal. Cultural similarity is the percentage of cultural dimensions that share the same trait values. An interaction between two neighbours happens with a probability proportional to their cultural similarity such that the focal actor adopts one trait value from the neighbouring actor.

Axelrod's model is set up to ensure a local convergence of culture, and yet, global polarisation may emerge from the model [829]. Such an outcome is more likely with fewer cultural dimensions, but lots of traits per dimension. In the context of immigrant integration, these results imply that without a concerted effort from the central government, social influence by natives may actually push immigrants aside rather than integrate them. Crises are likely to exacerbate the problem because they tend to narrow the number of cultural dimensions that pervade political discourse as exemplified in the introduction. The shrinkage in the number of relevant cultural dimensions is consistent with the mathematical formalism of phase transitions in the vicinity of tipping points, suggesting a way for quantitative surveys to measure how far a population is from the tipping point at which radical societal and political changes become probable.

Speaking of phase transitions, physicists have extended Axelrod's social-influence model and subjected it to intensive study in order to understand the model's dynamics [831, 832, 833, 834, 835]. Ref. [831], in particular, discusses a phase transition in Axelrod's model separating an ordered from a disordered phase. In the ordered phase, a dominant cultural region spans a large fraction of the whole system, whereas in the disordered phase, the system's state is fragmented into many cultural regions whose sizes are distributed in a non-trivial manner. The transition turns from continuous to discontinuous as the number of cultural dimensions increases. In relation to immigration policies, these results imply that seemingly similar conditions may lead to greatly

different outcomes; there could be a dominant cultural region enveloping most of society, or society could get shattered into many isolated cultural regions. It is questionable whether this latter outcome is compatible with modern-day nation states.

Unlike Axelrod's model in which local convergence is encouraged, there are models of social dynamics that favour divergent individuals [836, 837, 838]. A typical example is the seceder model [839, 840, 841]. This model leads to complex group formation such that, at random times, new groups split from old ones or existing groups go extinct. In medium-sized populations, the distance between two groups that are furthest apart tends to saturate, but in large-sized populations, this distance seems to increase linearly forever. The model thus mimics how subcultures pop in and out of existence [842], but may also offer insights into how immigrant communities develop while seeking to preserve cultural heritage and uniqueness. Conditions that lead to radicalisation, or the prevention thereof, could perhaps be better understood by drawing inspiration from the seceder model.

9.3. Populism and polarisation

Since its rise in the 1980s, populism has become a tool wielded by parties across the whole political spectrum [843]. In the context of migrations, however, it is right-wing populism that is of most interest. This particular type of populism is widely seen as pathological and pseudo-democratic in the sense that it is accompanied by a radically xenophobic and authoritarian political programme [844].

Examining the determinants of right-wing populism, Ref. [845] suggests that economic parameters play a smaller role than it is often assumed. An analysis of the results of the European Social Survey points in the direction that the electorates of the right-wing populist parties are more afraid of the negative influence of immigrants on a country's culture and heritage than on the country's economy. Data, in fact, suggests that low unemployment rates provide a fertile soil for the growth of right-wing populism [846, 847]. Ref. [848] furthermore finds that right-wing populist parties benefit from more crime, especially by linking crime to more immigration.

Analysing poll data from a group of EU countries affected by the recent migrant crisis, Ref. [535] finds that over the three-year period from 2014 to 2016 (i.e., in the midst of the migrant crisis), the percentage of right-wing populist voters in a given country depended on the prevalence of immigrants in this country's population and the total immigration inflow into the entire

EU. The latter was likely due to the perception that the EU functions as a supranational state in which a lack of inner borders means that ‘someone else’s problem’ can easily become ‘my problem’. When the annual immigrant inflow exceeded 0.4 % of a country’s population, it invariably led to an annual increase in the right-wing populist voters anywhere between 1 % and 5 %, implying that a prolonged large inflow could eventually cause right-wing populism to prevail.

Ref. [535] proceeds to mechanistically describe the rise of right-wing populism using a network-science model that accounts for the existence of tipping points in social dynamics. The model is constructed by placing a constant number of native ‘insider’ agents in a random network of social contacts. Immigrant ‘outsider’ agents subsequently enter the network. Each insider notices the percentage of outsiders in their neighbourhood and based on this percentage decides whether or not to turn to right-wing populism. Such a decision is based on local information, but non-local information can also affect decisions, as can misinformation. Three assumptions formalise these ideas:

1. *Global influences*, such as unfavourable socioeconomics and media reports, are assumed to induce a small negative bias in the decision making of any insider agent anywhere in the network. Direct contact with immigrants is unnecessary for anti-immigrant sentiments as evidenced by the BREXIT referendum in which low-immigrant areas mainly voted Leave [849].
2. *Opinion contagion*, complementing global influences, allows the seeds of populism to take root and turn into a full fledged right-wing populist movement. The contagion happens because of connections between insider agents. Namely, when an insider agent is surrounded by a critical number of right-wing populist-supporting neighbours, the agent’s decision making becomes negatively biased.
3. *Local influences*, specifically the perceived abundance of outsiders, is assumed to negatively bias the decision making of insider agents whose tolerance threshold has been exceeded by the number of outsider neighbours. For example, in local elections in Greece in November 2010, the far-right Golden Dawn party received only 5.3 % of the vote, but in some neighbourhoods of Athens with large immigrant communities the party won nearly 20 %.

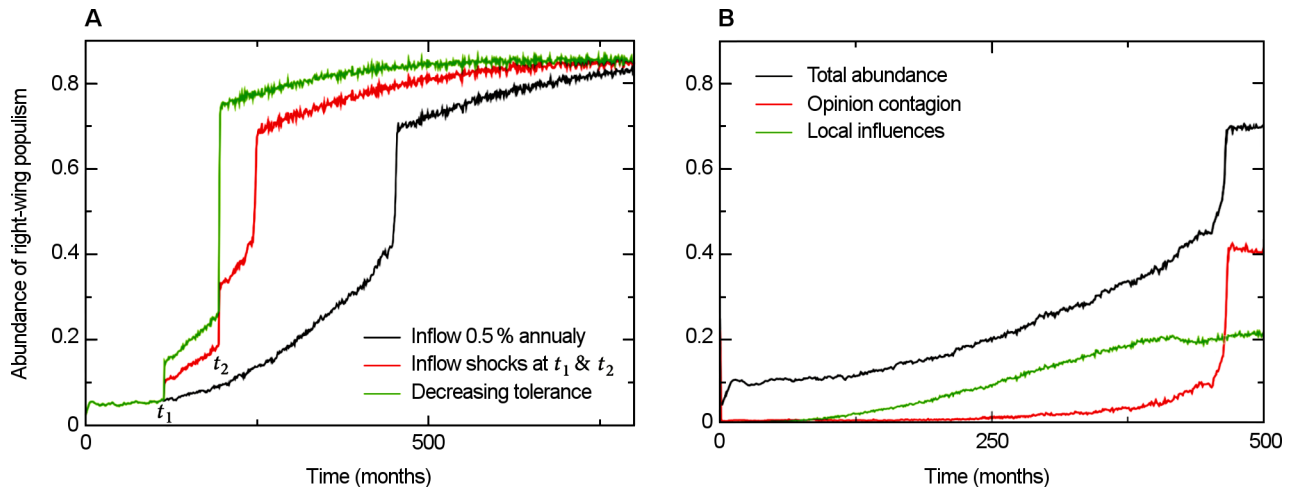


Figure 53: Rise of right-wing populism in a finitely tolerant population. **A**, Simple assumptions about the interactions of insider and outsider agents in a social network lead to a non-linear dynamics and discontinuous jumps in the abundance of right-wing populist supporters. **B**, Breakdown of the causes behind right-wing populism reveals that relatively far from the tipping point, the abundance of right-wing populist supporters responds to local influences. As the network approaches its tipping point, however, opinion contagion takes over and accelerates the transition to society dominated by right-wing populism. Source: Reprinted figure from Ref. [535] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

Simulation runs with an annual outsider inflow of 0.5% of the total population show that a tipping point starts manifesting itself as the fraction of outsiders approaches the tolerance threshold of insiders. The abundance of right-wing populist supporters increases non-linearly and eventually undergoes a sudden, discontinuous jump at about 37 years (450 months) into the simulation (black curve in Fig. 53A). The jump occurs much earlier if there are inflow shocks. Such shocks happen at times t_1 and t_2 , and cause inflows that are equivalent to about 5% of the total population. The closer the system to the tipping point, the effect of exactly the same shock becomes disproportionately larger (red curve in Fig. 53A).

Examining contributions to the rise of right-wing populism, it is evident that global influences can seed populist ideas here and there, but such ideas cannot be sustained without other processes. At first it is local influences that drive the increase in the abundance of right-wing populist supporters in the network (Fig. 53B). Opinion contagion remains a relatively small con-

tributor until the tipping point is approached (Fig. 53B). Near the tipping point, however, opinion contagion is explosive and overtakes local influences as the main source of right-wing populism. Thereafter right-wing populist supporters dominate society.

Right-wing populist policies are deeply divisive with a large potential to polarise societies. Similarly to how models of social influence helped us to gain insights into cultural integration in the preceding section, here we rely on models of opinion dynamics to gain insights into social polarisation. Ref. [850] is an early and influential work in this context, examining the dynamics of continuous opinions in well-mixed and lattice-structured populations. Continuous opinions nicely correspond to a variety of possible positions on the political spectrum.

The opinion-dynamics model in Ref. [850] has a very simple structure that in some aspects resembles the structure of previously discussed Axelrod's model of social influence. Each agent holds an opinion $x_i \in [0, 1]$ that is initially drawn at random from the uniform distribution. In simulations, a pair of agents meet either by chance (well-mixed case) or because they are neighbours in the social network (lattice-structured case). If their opinions are closer than a threshold d , $|x_i - x_j| < d$, the two agents' opinions approach one another at a convergence rate μ ; otherwise, the opinions remain unchanged. Just as two cultures interact only if they share some common points, so do two opinions interact only if they are close enough to begin with.

Simulation results reveal a critical role of the threshold d . Opinions converge by the very nature of the model, but consensus is observed only for large enough d values (Fig. 54A). For smaller d values the population polarises, first forming two distinct opinions whose gap cannot be closed (Fig. 54B), and then even more, with the approximate number of distinct opinions being $\lfloor 2d \rfloor$. In a lattice-structured population, the results are similar, although not without some additional interesting properties. For large enough d values, one percolating opinion dominates the lattice, while a small number of isolated opinions get randomly scattered around. As the d values decrease, it is still possible to observe one percolating opinion, but also a few sizeable clusters of distinct opinions. Interestingly, opinions within any single non-percolating cluster are quite similar but not entirely identical.

Interpreting the described results in the context of migrations, when less tolerant natives start resorting to right-wing populism, many of the right-wing populist policies are unpalatable to more tolerant natives, thus creating




Figure 54: Opinion dynamics may lead to a consensus or polarise society. **A**, Assuming that individuals tolerate opinions that differ from their own opinion by at most a distance d , sufficiently large d values lead to consensus. **B**, As the d value decreases, society gets polarised between two competing opinions. Even smaller d values cause more than two competing opinions to persevere. Source: Reprinted figure from Ref. [850].

a rift in the population. If communication across the rift is very limited, opinion dynamics suggests that society is likely to polarise. Aside from immigration, evidence shows that a similar, catalytic role is played by partisanship, religious orientation, and even geographical differences [851, 852]. Polarisation, furthermore, leads to more political participation [851], hinting that once a rift forms, each sub-population doubles down on its own opinion. Evidence that rifts across which limited communication occur is found in the phenomenon of echo chambers, that is, groups of like-minded individuals who subscribe to, and mutually reinforce, a certain narrative [853, 854, 855, 856, 857]. Of note here is that the extent to which online echo chambers impact society is a matter of debate [858].

Similar to Axelrod’s model of social influence, the opinion-dynamics model of Deffuant et al. [850] has attracted a lot of attention among physicists [859, 860, 861, 862, 863]. In adaptive networks [864], for instance, agents can sever links with those agents who harbour very different opinions, while preferen-

tially linking with those agents who harbour similar opinions. Such adaptability reduces the chances of reaching both the consensus state and any of the highly fragmented states, thus in effect promoting polarisation [865]. Extending the opinion-dynamics model even further, a recent study aims at explaining the above-mentioned phenomenon of echo chambers and polarisation on social media [866]. This same body of research also encompasses voting phenomena and the emergence of language. The voter model is concerned with capturing opinion dynamics under a strong influence of an individual's social context [867, 868, 869], including statistical regularities of real-world electoral processes [870, 871]. The theory of language emergence is concerned with communication strategies employed by individual speakers, and how the systematic use of such strategies leads to a consensus by which one word conveys the same idea to everyone [872, 873, 874, 875, 876, 877]. All this goes to show that the research on opinion dynamics has an enormous scope and breadth, with many important contributions yet to come.

9.4. Future outlook

Although it could be argued that the geopolitical situation leading to the migrant crisis that hit Europe in recent years is a one-off event, neither the EU nor other destination countries have the luxury of ignoring large-scale migrations in the near- to mid-term future. Climate change, for example, is expected to displace millions of people over the next few decades [814]. Worse yet, the recent crisis taught us that large-scale migrations can be accompanied with immeasurable suffering and the loss of human life. For these reasons, we see the need for an interdisciplinary research agenda whose main goal would be attaining a quantitative understanding of migrations and the underlying social dilemmas. The necessary scientific tools that could aid sociological and economic research are available in the form of network science (collective phenomena), statistical physics (tipping points and phase transitions), evolutionary game theory (cooperative Nash equilibria), and others. We furthermore believe that modern approaches should be data driven, and therefore rely on the methods of Bayesian statistics, econometrics, and machine learning.

Fast and slow migrations. Current research indicates that serious imbalances between (i) the inflow of migrants, (ii) the willingness of the native population to accept them, and (iii) migrant integration causes a knee-jerk political response by which right-wing populists are voted into offices. Subsequently,

discriminatory laws and regulations targeting migrants or underrepresented minorities are often proposed and even enacted. It thus becomes essential to consider the role of speed at which processes in social dynamics unravel. If the inflow of migrants is moderate, the native population should feel secure and adapt, but if the inflow is extremely high, the natives may perceive migrants as a threat and—in a sort of Popperian twist—respond antagonistically. These two limits suggest that tolerance is a relative category depending on the magnitude of immigrant inflow, which is an effect that must be included in a genuine model of large-scale migrations. Otherwise, rising right-wing populism may lead to a demise of values such as free thought, liberty, democracy, human rights, and even peace.

Migration dilemma. Countries today face a migration dilemma that has two complementary dimensions. The first dimension is faced by countries with a positive net inflow of migrants (i.e., immigrants). Such countries, comprising mostly ‘old democracies’, must decide whether to decline or accept immigrants, and in the latter case, decide the rate of acceptance before the integration capacity is fully utilised. Therein lies the dilemma for old democracies. The second dimension is faced by countries with a positive net outflow of migrants (i.e., emigrants), usually economic ones, who seek better life elsewhere. Such countries, comprising mostly ‘new democracies’, must decide whether to replenish their populations with immigrants or face socio-economic instabilities in the future. Therein lies the dilemma for new democracies. Irrespective of how individual countries choose to resolve the dilemma, it is important to recognise that the consequences will likely be non-local because of strong socio-economic interdependencies in a globalised world. This finally returns us to network science, statistical physics, and evolutionary game theory as theoretical frameworks equipped to handle interdependent complex systems [266], tipping points [878], and incentives for human behaviour [233]. Intertwining these frameworks with sociological and economic knowledge might provide us with a detailed enough picture of how to resolve the migration dilemma for the benefit of all.

10. Contagion phenomena

At the time of writing, the Covid-19 pandemic has been raging around the world for nearly two years, while the cause of this pandemic, the SARS-CoV-2 virus, has infected over 200 million people. At this scale, the pandemic is

of course a global public health issue of which general population is acutely aware. Viral outbreaks, however, oftentimes fall short of such widespread awareness despite their recurrent occurrences and deadly potential. To understand just how frequent and dangerous viral outbreaks are, it is illustrative to consider the following examples. Influenza A(H1N1) outbreak started in Mexico in 2009 and subsequently reached 214 countries and regions, causing 18,500 deaths [879]. Avian influenza A(H7N9) outbreak started in mainland China in 2013, giving rise to only 419 cases because of no sustained human-human transmission, but still causing 127 deaths due to high virulence in humans [880]. Ebola virus disease emerged in West Africa in 2014 and went on to infect more than 25,000 people, resulting in over 11,000 known fatalities although the true case fatality rate is suspected to be above 70 % [881]. Finally, the MERS-CoV outbreak started in the Middle East in 2012 and proceeded to cause over 1,350 human infections with a death toll of more than 500 people from 26 countries [882].

For nearly 100 years, the bread and butter of research into disease outbreaks have been compartmental epidemiological models. It was Kermack and McKendrick who in 1927 put forth a mathematical framework for compartmental models in epidemiology [883]. Furthermore, the work of Reed and Frost from about the same time, but exposed later by others [884], had been the first to introduce a chain-binomial model with the recognisable susceptible-infectious-recovered (SIR) structure. In SIR models, a population exposed to an active pathogen is divided into three compartments. Healthy individuals are susceptible to the pathogen (S), infectious individuals transmit the pathogen (I), while recovered individuals no longer respond to the pathogen due to acquired immunity (R). The simplest SIR model is therefore given by

$$\frac{dS}{dt} = -\beta SI, \quad (87a)$$

$$\frac{dI}{dt} = \beta SI - \gamma I, \quad (87b)$$

$$\frac{dR}{dt} = \gamma I, \quad (87c)$$

where β and γ are infection and recovery rates, respectively. The pathogen cannot spread if $\frac{dI}{dt} \leq 0$, which leads to the condition $R_0 = \frac{\beta S(0)}{\gamma} \leq 1$. The quantity R_0 is called the basic reproductive number, and stands for the expected number of cases caused by a single case in a completely susceptible

population. Once the epidemic takes off, its decline is subject to the condition $R_e = \frac{\beta S(t)}{\gamma}$, in which case we are talking about an effective reproductive number because the population is no longer completely susceptible. Among the common extensions of the model is to add a population growth rate Λ , as well as the different mortality rates for each compartment d_i , $i \in \{S, I, R\}$. The model then takes the shape

$$\frac{dS}{dt} = \Lambda - \beta SI - d_S S, \quad (88a)$$

$$\frac{dI}{dt} = \beta SI - \gamma I - d_I I, \quad (88b)$$

$$\frac{dR}{dt} = \gamma I - d_R R. \quad (88c)$$

Another straightforward extension of the SIR model is to consider a network of populations (e.g., cities). The model then becomes

$$\frac{dS_i}{dt} = \Lambda_i - \beta_i S_i I_i - d_i^S S_i + \sum_{i=1}^n a_{ij} S_j - S_i \sum_{i=1}^n a_{ji}, \quad (89a)$$

$$\frac{dI_i}{dt} = \beta_i S_i I_i - \gamma_i I_i - d_i^I I_i + \sum_{i=1}^n b_{ij} I_j - I_i \sum_{i=1}^n b_{ji}, \quad (89b)$$

$$\frac{dR_i}{dt} = \gamma_i I_i - d_i^R R_i + \sum_{i=1}^n c_{ij} R_j - R_i \sum_{i=1}^n c_{ji}, \quad (89c)$$

where the matrices $A = [a_{ij}]$, $B = [b_{ij}]$, and $C = [c_{ij}]$ quantify mobility rates between populations i and j [885]. Models of this type are often referred to as metapopulation epidemiological models. Since the 1980s, metapopulation models built around data on the worldwide air-transportation network [886] have become one of the main tools for studying the global spread of emerging epidemics. This is also where our focus lies in the subsequent sections of this chapter.

10.1. Data-driven metapopulation models

Metapopulation epidemiological models are built as a network of populations such that transmission dynamics occurs on two scales. The smaller scale describes the local disease spread within each population, while the larger scale describes the disease spread between between populations due to the movements of infectious individuals.

Disease spread within a population. Let us consider a disease that can be described by an SIR compartmental model, such as pandemic influenza or measles. Susceptible individuals (S) become infectious (I) at a rate β after encountering an infectious individual, whereas each infectious individual becomes recovered or dies at a rate μ . Assuming that the contacts between susceptible and infectious individuals are frequency dependent, we can model transitions in population i from susceptible to infectious individuals and from infectious to recovered individuals over the unit time period Δt using the following equations

$$S_i(t + \Delta t) = S_i(t) - \Delta S_i(t), \quad (90a)$$

$$I_i(t + \Delta t) = I_i(t) + \Delta S_i(t) - \Delta I_i(t), \quad (90b)$$

$$R_i(t + \Delta t) = R_i(t) + \Delta I_i(t) \quad (90c)$$

$$\Delta S_i(t) \sim \mathcal{B}(S_i(t), \beta \frac{I_i(t)}{N_i} \Delta t) \quad (90d)$$

$$\Delta I_i(t) \sim \mathcal{B}(I_i(t), \mu \Delta t) \quad (90e)$$

where $\mathcal{B}(n, p)$ denotes a binomial random variable with the parameters n for the number of trials and p for the probability of success, whereas $N_i = S_i(t) + I_i(t) + R_i(t)$ denotes the size of the i th population. For pathogens with a relatively short doubling time (e.g., SARS-CoV-2, MERS-CoV, influenza A(H1N1), Ebola, measles, etc.), it is appropriate to set a short unit time period (e.g., $\Delta t = 0.05$ d).

Disease spread between populations. In metapopulation models, individuals are imagined to jump between populations using a transportation network. The simplest scenario is to consider that individual movement is stochastic, which approximates well the international spread of infectious diseases [887, 888, 889, 890]. To account for more complex patterns of human mobility, metapopulation models have also been extended to address the memory effects of individual mobility (e.g., daily commuting) [891, 892, 893], differential social mixing patterns due to socioeconomic stratification [894], and age-related factors [895]. Using as an example the simplest metapopulation epidemiological model with G populations, the spread of epidemics from each population i to downstream populations that are directly connected to

population i can be described using

$$X_i(t) = \{X_{i1}(t), \dots, X_{iG}(t)\} \sim \mathcal{M}(S_i(t), w_{i1}\Delta t, \dots, w_{iG}\Delta t), \quad (91a)$$

$$Y_i(t) = \{Y_{i1}(t), \dots, Y_{iG}(t)\} \sim \mathcal{M}(I_i(t), w_{i1}\Delta t, \dots, w_{iG}\Delta t), \quad (91b)$$

$$Z_i(t) = \{Z_{i1}(t), \dots, Z_{iG}(t)\} \sim \mathcal{M}(R_i(t), w_{i1}\Delta t, \dots, w_{iG}\Delta t), \quad (91c)$$

where \mathcal{M} denotes a multinomial random variable, w_{ij} denotes the mobility rate between populations i and j , $X_{ij}(t)$, $Y_{ij}(t)$, and $Z_{ij}(t)$ respectively denote the number of susceptible, infectious, and recovered individuals who travel from population i to population j between the times t and $t + \Delta t$.

The need for data. To build metapopulation epidemiological models for studying real-world epidemics, it is essential to parameterise the connectivity of the underlying transportation network, as well as population flows along this network. This has been made possible by recent advancements in digital data collection and storage (see Chapter 7). Examples of suitable data sources include:

- *Official Aviation Guide*, a data-subscription service covering the worldwide flight booking database that has previously been used for building state-of-the-art global epidemic simulators (e.g., GLEAM) [893].
- *mobile service providers* (e.g., the Telenor Group or Orange) whose data-subscription services offer access to anonymous call detail records (CDRs), which allow the quantification of aggregated population movements between cities [896, 897].
- *search-engine and social-media companies* (e.g., Google, Baidu, Facebook, Tencent, etc.) whose open-access data portals or data-subscription services offer population-movement information derived from mobile location-based services (LBS) [898, 899].

Compared to the CDR data, the LBS data allows the stratification of population flows by points of interests or modes of transportation. For example, Google’s Covid-19 Community Mobility Reports provide six data-streams called ‘grocery and pharmacy’, ‘parks’, ‘transit stations’, ‘retail and recreation’, ‘residential’, and ‘workplaces’ [900]. Tencent’s migration-data portal separates travel by aeroplanes, railways, or highways [901]. Overall, modern sources of digital data usable in epidemiology have become so rich that the term ‘digital epidemiology’ is now being used to describe the extent to which

epidemiologists have come to rely on such sources. This is the topic of our next section.

10.2. Digital epidemiology

Definition. Relative to traditional epidemiology that uses data generated by the public health system to understand the incidence, distribution, and possible control of diseases and other health-related factors, digital epidemiology uses digital data generated elsewhere [902]. Among the main reasons for this is that data collected by professionals in clinics, hospitals, and laboratories is accurate but very costly in terms of logistics, time, labour, and materials. Tapping into alternative and cost-efficient sources therefore seems like a natural way to move forward [903].

Internet and mobile-phone uses result in billions of digital-communication records that document health-related behaviours such as symptom reports or attitudes towards vaccines. For example, Google searches provide estimates of influenza activity in a near real-time manner [530], whereas mobile-phone data tracks population movements during disease outbreaks [904]. Such capabilities have led to the applications of digital epidemiology in surveilling and predicting the spread of air-borne and vector-borne viruses, parasites, and other pathogens.

Data sources. Google Trends is a free service providing normalised trends in search activity depending on user-specified geographical regions and time frames. Refs. [905, 530] found that the relative frequency of certain queries about influenza is closely correlated with influenza cases reported by public health agencies in the United States. These pioneering works thus opened the door to using search queries to detect near real-time influenza epidemics.

The original algorithm in Refs. [905, 530] has later been shown to suffer from several major limitations that lead to inaccurate estimates [906, 907]. The algorithm is static and fails to account for time-series properties such as the seasonality of influenza activity, whereas aggregating multiple query terms into a single variable ignores changes in internet-search behaviour over time. In response, Ref. [908] proposed the AutoRegression with Google search data (ARGO) model to address the above-said shortcomings. Based on the ARGO model, Ref. [909] further improved influenza prediction by incorporating spatial and temporal synchronicities seen in historical flu activity at the state-level in the US.

In addition to success stories from developed countries such as the data-rich United States, Google Trends has proven valuable in developing countries with poorer data environments for infectious disease nowcasting and forecasting. For example, Ref. [910] has demonstrated the power of the ARGO model to improve influenza predictions for 8 different Latin-American countries: Argentina, Bolivia, Brazil, Chile, Mexico, Paraguay, Peru, and Uruguay.

In addition to monitoring influenza-like illnesses, internet search queries have also been applied to vector-borne diseases such as dengue, malaria, and Chagas disease in tropical and temperate low- to middle-income countries [911, 912, 913, 914, 915]. Because of an increasing internet-access availability, but relatively limited traditional epidemiological data-collection capabilities [916], internet-based surveillance methods have the potential to greatly complement the work of public-health agencies on preparedness against vector-borne diseases. Two independent studies [911, 912] have shown that web search query data has a high correlation ranging from 0.82 to 0.99 with dengue activity in Bolivia, Brazil, India, Indonesia, and Singapore. Afterwards, Google developed a prediction tool called Google Dengue Trends to provide timely information to public health agencies from Mexico [917], Taiwan [918], Venezuela [919], and the Philippines [920].

Twitter is another big-data source of interest to public-health researchers because of the real-time nature of content, precise geotagged locations, and publicly available information. Tweet-based disease surveillance and prediction exploits time-series trends in health-related keyword volumes as a predictor variable in regression models, which is conceptually similar to how Google Trends is used. However, because most tweets employ everyday language to describe a combination of symptoms rather than a diagnosis, there is a need for natural-language processing in order to identify relevant information. Ref. [921] found a large correlation coefficient between the trends of influenza mentions on Twitter and influenza-like illnesses identified by traditional surveillance systems in the US. Ref. [922] further recorded a high prediction accuracy (85 %) in relation to the weekly change of influenza prevalence on both national and city scales (e.g., in the New York city).

When it comes to vector-borne diseases, Twitter has proven useful in surveying and predicting dengue in Brazil on national, regional, and city scales [923]. The best performance has been observed on the national scale, on which the correlation coefficient between predictions and dengue activity was as large as 0.98.

Twitter is also an important source of data to assess health-related be-

haviours such as concerns about disease outbreaks and attitudes to public-health measures [924, 925]. Ref. [926] relied on publicly available Twitter data to resolve the spatio-temporal sentiment towards a new influenza A (H1N1) vaccine in 2009. The study identified a strong correlation between the online sentiment and estimated vaccination rates by region. Subsequent data-driven simulations of disease spread found that the clusters of a negative vaccine sentiment tend to coincide with an increased likelihood of disease outbreaks. Ref. [927] additionally found that sentiments themselves are contagious. Negative sentiments in particular are more contagious than positive ones.

To exhaust the potential of Twitter data, the following two applications must also be considered. Tweets provide an opportunity to discuss medications, which perhaps could be used to detect drug-related adverse events and thus improve pharmacovigilance. Ref. [928] showcases an analysis of 2 billion tweets in search of adverse events related to 5 different cancer drugs. The analysis identified 239 potential drug users. Each potential case was then examined by two experts, which led to 72 definite positives. From these positives, 27 drug-related adverse event were detected, thus providing a proof-of-concept solution towards improving pharmacovigilance based on Twitter data. The other application arises from the fact that tweets are geo-referenced. This implies the possibility to capture human-movement patterns in order to track and control infectious diseases. Twitter stores geographic coordinates that offer insights into movements on various temporal (from daily onward) and spatial (from local to national to international) scales [929, 930].

Facebook is among the most visited websites in the world, but has not been used as much in public-health contexts because of limited data access in the past. Facebook Data for Good is a recent project aimed at broadening access for social-welfare purposes [931]. One of the major advantages of Facebook data is the information on social connections. Albeit these connections are established in an online environment, there is a strong correlation with the geography of health-related activities. Ref. [932] thus found that Covid-19 tends to spread between regions with more social-network connections as indicated by Facebook. This showcases that data from online social networks could be used to forecast the spread of air-borne diseases based on proximity indicators derived from digital interactions.

Integrating the information on human movements into epidemiological models leads to insights into the disease spread, as well as the optimal resource allocation to contain the spread. Ref. [933] is an attempt to do so

using Facebook movement data while evaluating the economic consequences of alternative lockdown-lifting scenarios in various Italian districts. The results show that there is a tradeoff between disease transmission and worker mobility such that a given economic loss on the national scale induces heterogeneous regional losses. Furthermore, humanitarian organisations need to know where to allocate resources to help people who are most affected by a disease outbreak or other disasters. Ref. [934] shows how aggregating Facebook usage in areas impacted by such events can be used to produce disaster maps outlining population evacuations and long-term displacements.

Ensemble estimates. Despite the apparent success of relying on digital-data sources for surveilling and predicting contagions, this methodology has been criticised and concerns have been raised [935, 936, 937, 906, 907]. Ensemble models have been developed in response to such criticisms, following the idea that combining multiple digital-data sources circumvents the weaknesses that any individual source might have.

Ref. [938] outlines an ensemble machine-learning model to predict influenza activity in the US by leveraging Google Trends, Twitter data, Flu Near You [939], and the CDC data on influenza-like illnesses. The results demonstrate that combining information from multiple data sources improves real-time predictions up to four weeks ahead. Encouraging results have also been obtained in middle-income countries from Latin America where available data is scarcer [940]. Similarly, Ref. [941] combines the information on Zika virus disease from Google searches, Twitter, HealthMap [942], and suspected cases during the 2015-2016 Latin American outbreak to predict new weekly cases up to three weeks ahead.

Refining spatial resolution decreases the correlation between predictions by digital-surveillance systems and estimates by public-health systems [943]. Ensemble models are a promising way forward in this context. Ref. [944], for example, combines official health reports, internet searches for Covid-19 on Baidu, news media, and results from an agent-based epidemiological model to produce accurate forecasts two days ahead on the provincial scale in China. Another similar example is Ref. [945] which combines Google searches, Twitter data, electronic health records, and Flu Near You [939] to predict influenza outbreaks in the Boston metropolitan area one week ahead (Fig. 55).

This brief overview by no means exhausts all possible sources of digital data that can be used in epidemiology. What is more, the variety of such

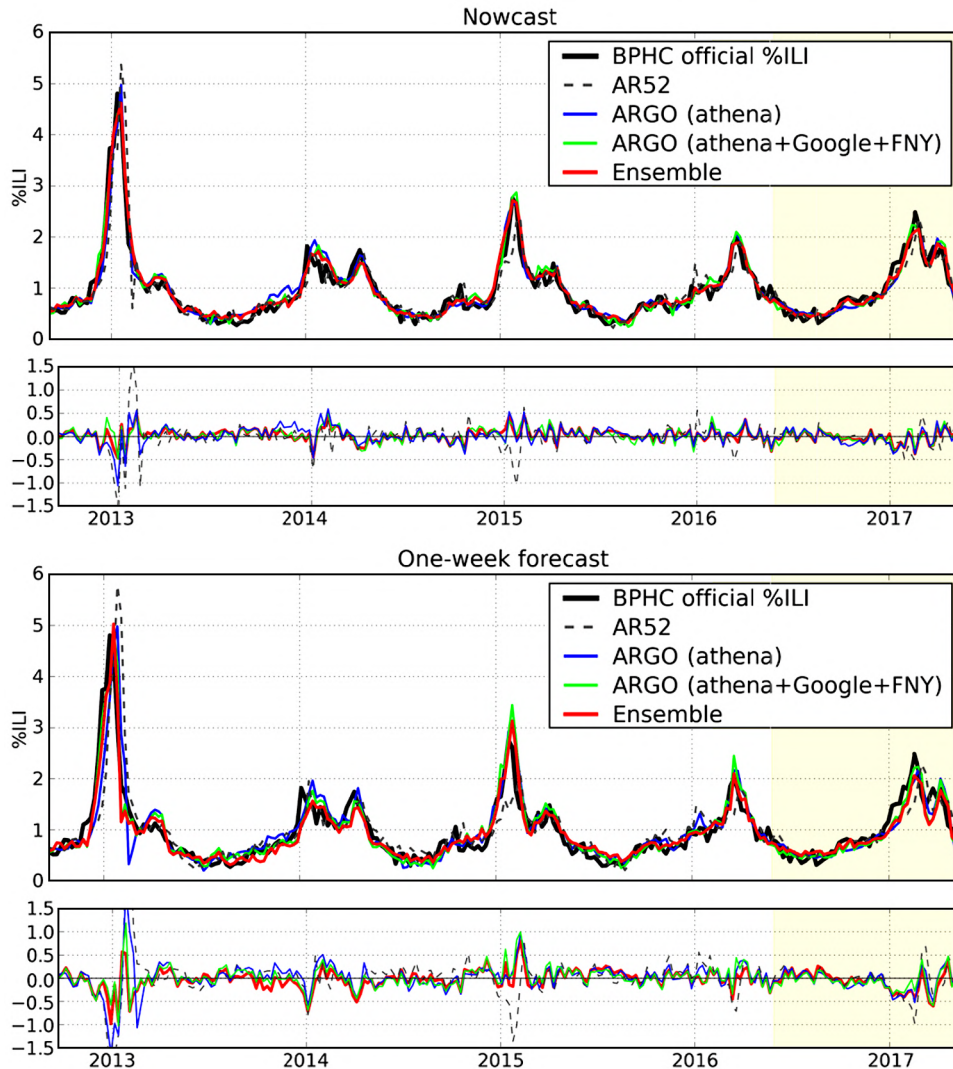


Figure 55: Nowcasts and one-week forecasts (with errors) of influenza-like illnesses in the Boston metropolitan area from September 2012 to May 2017. BPHC stands for Boston Public Health Commission. AR52 stands for an autoregressive baseline model using 52 weeks of past data to make predictions. ARGO stands for autoregression with general online information with the sources of online information being athenahealth (athena), Google Trends (Google), and Flu Near You (FNY).
 Source: Reprinted figure from Ref. [945] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

sources is bound to increase in the short- to mid-term future. A key development, however, will be to couple them with models that offer mechanistic insights into epidemics. This, of course, includes state-of-the-art metapopulation models discussed in detail in this chapter.

10.3. Analytical results from metapopulation models

From the viewpoint of statistical physics, the epidemic threshold is an important concept arising inspired by the studies of critical phenomena and phase transitions [946]. The epidemic threshold delineates where in the parameter space an epidemic wanes and where in the parameter space the epidemic intensifies. Interestingly, studying the spread of computer viruses on the Internet has shown that the epidemic threshold can be negligibly small if the node-degree distribution of a network is strongly heterogeneous, as is the case for scale-free networks [947, 948]. This was a surprising finding at the time, motivating various searches for novel public-health policies. For example, subsequent attempts to formulate effective vaccination strategies in networks have largely focused on vaccinating hub nodes, that is, those nodes that possess the most connections [949].

The research on epidemic threshold has eventually been extended to metapopulation epidemiological models. Refs. [950, 951], for example, describe a mean-field derivation leading to a closed-form expression for the epidemic threshold of a multiple-population network in which individuals randomly move between populations. The threshold was found to decrease as the network becomes more node-degree heterogeneous. Refs. [952, 953] extend this line of work to account for recurrent mobility patterns.

A concept of interest in addition to the epidemic threshold is the epidemic arrival time (EAT). After a disease outbreak in a city of origin, for example, Wuhan in case of the Covid-19 pandemic, the disease can spread to other cities through the travel of individuals. The EAT for each downstream city j is the time at which the first infected case is imported into this city. The EAT measures the spreading velocity of the disease and encodes relatively reliable information for inferring key epidemiological parameters in the early phases of novel epidemics [954, 890, 955, 956].

Although several seminal studies [957, 958, 959] have explored the potential for developing a simple summary statistics to approximate the EAT, a general analytical framework leading to a closed-form expression for the probability distribution of the EAT has remained elusive. To fill this knowledge gap, Ref. [890] derives the probability distribution of the EAT in three

metapopulation models with increasingly complex network structure: (i) the simplest two-population model, (ii) the shortest-path tree of the worldwide air-transportation network, and (iii) the whole worldwide air-transportation network.

Two-population model. The simplest problem for which the EAT distribution can be found analytically is that of two populations. An infectious disease is assumed to originate from population i which is in turn connected to population j (Fig. 56A). This situation corresponds to initial stages of an epidemic when new infections emerge in the origin population, while all the other populations are aggregated into a single population momentarily unaffected by the disease [889]. Two key mathematical assumptions [890, 957, 958] made at this point are:

1. Exportation of infections from population i to j is a non-homogeneous Poisson process [960] with the intensity function (i.e., the expected number of infections exported from population i to population j at time t) given by $w_{ij}I_i(t)$, where $I_i(t)$ is the number of infectious people (i.e., disease prevalence) in population i at time t , and w_{ij} is the per-capita mobility rate from population i to population j .
2. After the new epidemic establishes itself in origin population i , the first few exportations from population i to population j occur while disease prevalence is still growing exponentially in the origin, that is, $I_i(t) = s_i \exp(\lambda_i t)$, where s_i is the initial seed size in origin population i at time $t = 0$, and λ_i is the local epidemic growth rate.

With the above-stated assumptions, the n th EAT in population j , denoted T_{ij}^n , is a random variable whose probability density function can be expressed in closed form

$$f_n(t|\lambda_i, \alpha_{ij}) = \left(\frac{e^{\lambda_i t} - 1}{\lambda_i} \right)^{n-1} \frac{\alpha_{ij}^n}{(n-1)!} \exp \left(\lambda_i t - \alpha_{ij} \frac{e^{\lambda_i t} - 1}{\lambda_i} \right), \quad (92)$$

where $\alpha_{ij} = s_i w_{ij}$ is the adjusted mobility rate. The last expression can be validated numerically. To this end, analytical and simulated EATs were compared over a wide range of epidemic scenarios (Fig. 56). Among others, doubling and generation times varied between 3 and 30 days, which was enough to cover many present-day infectious diseases. Covid-19 has a doubling time of 5-7 days, whereas Ebola has a doubling time of more than 20 days.




Figure 56: Two-population epidemiological model. **A**, Model schematics showing the situation in which origin population i connects only to population j . **B–D**, Q-Q plots showing the analytical and simulated quantiles of the random variables T_{ij}^1 , T_{ij}^5 , and T_{ij}^{10} . Insets show the corresponding histograms of the percent error in $E[T_{ij}^n]$. Simulations entail 100 epidemic scenarios sampled using the Latin-hypercube sampling from the following parameter space. Doubling and generation times both ranged between 3 and 30 days, the seed size s_i ranged between 1 and 100, the mobility rate w_{ij} ranged between 10^{-6} and 10^{-3} , and the population size N_i ranged between 0.1 and 10 million. The latter two parameters were respectively chosen according to the Official Airline Guide (OAG) air-traffic data and the Gridded Population of the World dataset (Version 4). Simulated quantiles for each of these 100 scenarios were computed from 10,000 stochastic realisations. Points on the diagonal indicate that analytically calculated and numerically simulated arrival-time quantiles are equal. Blue and yellow points distinguish scenarios in which $P(X_{ij} \geq n) = 1$ from those in which $P(X_{ij} \geq n) < 1$, where X_{ij} is the number of exportations. Source: Reprinted figure from Ref. [890] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

Eq. (92) can be used to derive a number of corollaries:

1. Exportation of the first n infections is a non-homogeneous Poisson process with the intensity function $\alpha_{ij} \exp(\lambda_i t)$.
2. The cumulative distribution function of the n th EAT is

$$F_n(t|\lambda_i, \alpha_{ij}) = \Gamma \left[n, \frac{\alpha_{ij}}{\lambda_i} (e^{\lambda_i t} - 1) \right], \quad (93)$$

where Γ denotes the lower incomplete Gamma function.

3. The expected first EAT is

$$E[T_{ij}^1] = \frac{1}{\lambda_i} \exp \left(\frac{\alpha_{ij}}{\lambda_i} \right) \mathbf{E}_1 \left(\frac{\alpha_{ij}}{\lambda_i} \right), \quad (94)$$

where $\mathbf{E}_m(x) = x^{m-1} \int_x^\infty \mu^{-m} e^{-\mu} d\mu$ is the exponential integral.

4. If $\alpha_{ij} \ll \lambda_i$ and γ denotes the Euler-Mascheroni constant, then the expected first EAT can be approximated using

$$E[T_{ij}^1] = \frac{1}{\lambda_i} \left[\ln \left(\frac{\lambda_i}{\alpha_{ij}} \right) - \gamma \right], \quad (95)$$

which is equivalent to the first-EAT statistic in Ref. [957].

5. The expected n th EAT is

$$E[T_{ij}^n] = \frac{1}{\lambda_i} \exp \left(\frac{\alpha_{ij}}{\lambda_i} \right) \sum_{m=1}^n \mathbf{E}_m \left(\frac{\alpha_{ij}}{\lambda_i} \right). \quad (96)$$

6. For any positive integers m and n , $m < n$, the probability density function of $T_{ij}^n - T_{ij}^m$ conditional on T_{ij}^m is

$$f_{n-m}(t | \lambda_i, \alpha_{ij} e^{\lambda_i T_{ij}^m}), \quad (97)$$

which can be reinterpreted as the probability density function of the $(n - m)$ th EAT with the seed size $s_i \exp(\lambda_i T_{ij}^m)$. Using this relation recursively, we deduce that the joint probability density function of $T_{ij}^1 = t_1, \dots, T_{ij}^n = t_n$ is

$$\prod_{m=1}^n f_1(t_m | \lambda_i, \alpha_{ij} e^{\lambda_i t_{m-1}}), \quad (98)$$

for all $0 = t_0 < t_1 < t_2 < \dots < t_{n-1} < t_n$.

7. The expected $(n - 1)$ th EAT given an epidemic that starts at time T_{ij}^1 with the seed size $s_i \exp(\lambda_i T_{ij}^1)$ is

$$E[T_{ij}^{n-1} | T_{ij}^1] = T_{ij}^1 + \frac{1}{\lambda_i} \exp \left(\frac{\alpha_{ij}}{\lambda_i} e^{\lambda_i T_{ij}^1} \right) \sum_{m=1}^{n-1} \mathbf{E}_m \left(\frac{\alpha_{ij}}{\lambda_i} e^{\lambda_i T_{ij}^1} \right). \quad (99)$$

These corollaries will prove useful in extending the two-population model to the analysis of the epidemic arrival process in the shortest-path tree of the worldwide air-transportation network, and the whole worldwide air-transportation network.

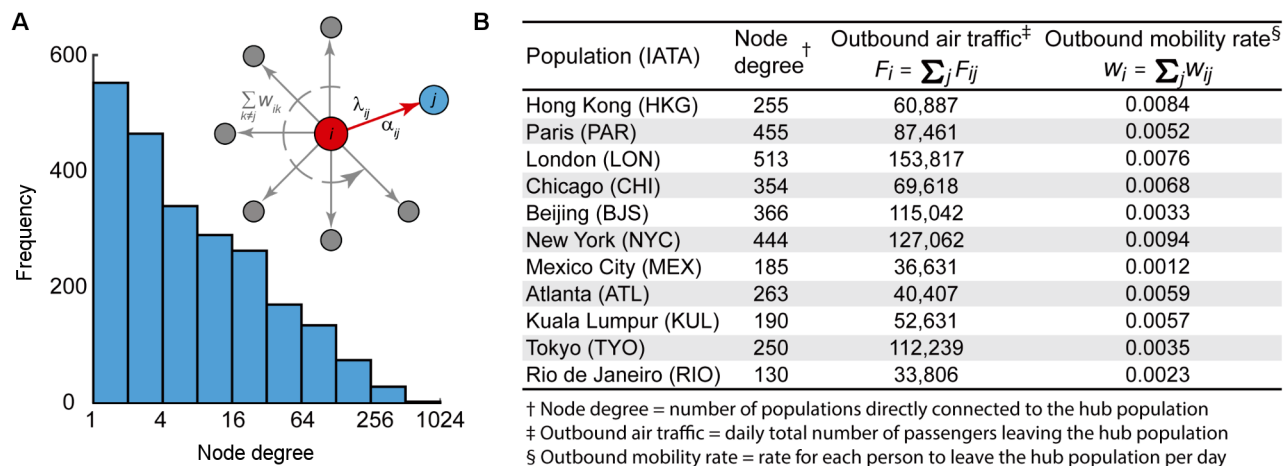


Figure 57: Network properties of hub populations. **A**, Histogram shows the distribution of node degrees for all populations in the worldwide air-transportation network. The node degree of a given population equals the number of populations to which this population is directly connected. The inset illustrates that travel-hub population i is connected to multiple populations, including population j . **B**, List of several major hub cities situated all over the world. Shown are the node degree, the daily outbound-traffic volume, and the daily outbound per-capita mobility rate.

Modelling the shortest-path tree. A dominant sub-network of the worldwide air-transportation network is its shortest-path tree. In this sub-network, each downstream population is connected to the epidemic origin via only one path. Ref. [959] suggests that an emerging epidemic spreads from the origin population to other populations mainly through the shortest-path tree, that is, the infrastructure of the shortest-path tree drives the global spread of the disease across the worldwide air-transportation network. Ref. [890] demonstrates that the n th EAT T_{ik}^n from origin population i to any population k in the shortest-path tree is accurately characterised by Eq. (92), but the local epidemic growth rate λ_i and the adjusted mobility rate α_{ij} need to be re-parameterised to account for what is called *the hub effect* (Fig. 57) and *the continuous-seeding effect* (Fig. 58).

Travel hubs such as Hong Kong, London, and Paris are characterised by direct flights to many locations. In network-science terminology, the node degree of travel hubs is much larger than unity. This creates many opportunities for infection exportation, perhaps to the point that the local growth of the disease prevalence, $I_i(t)$, is noticeably reduced. If indeed a noticeable




Figure 58: Continuous-seeding effect. **A**, Schematic of the epidemic arrival process (mathematically, a non-homogenous Poisson process) over an acyclic path connecting origin population i to population k via population j (i.e., $\psi : i \rightarrow j \rightarrow k$). **B**, In this example, the epidemic arrives in population k after population j has imported three infections from the origin, that is, $T_{ij}^3 < T_{ik}^1 < T_{ij}^4$. In the absence of continuous seeding adjustment, infection trees spawned by the second and subsequent importations in population j are ignored. Source: Reprinted figure from Ref. [890] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

proportion of infections travel outward as the epidemic unfolds, the local epidemic growth rate, λ_i , needs an adjustment.

Suppose that hub population i is directly connected to two or more populations, one of which is population j (Fig. 57A). In the shortest-path tree, all infectious individuals who disperse from population i to populations other than j no longer contribute to disease exportations to population j . A consequence is that the probability density function of the n th EAT is still given by Eq. (92), but the local epidemic growth rate in hub population i from the perspective of population j must be adjusted to

$$\lambda_{ij} = \lambda_i - \sum_{k \neq j} w_{ik}. \quad (100)$$

The random variable T_{ij}^n therefore has the probability density function $f_n(t|\lambda_{ij}, \alpha_{ij})$, implying that the disease prevalence in hub population i grows exponentially at the effective growth rate λ_{ij} , while the number of exported infections from population i to population j at time t remains the same as before, that is, $w_{ij}I_i(t)$. The need for the described adjustment can be numerically validated in a similar manner as the two-population model (Fig. 59A).




Figure 59: Numerical validation of analytics for the shortest-path tree of the worldwide air-transportation network. Shown are the Q-Q plots comparing the analytical and simulated quantiles of various EATs for downstream populations in the shortest-path tree. Insets show the corresponding histograms of the percent error in the expected EATs. The origin of the epidemic was assumed to be in Hong Kong. The same 100 epidemic scenarios as in Fig. 56 were used. The symbol D_c stands for the set of all populations that are separated by c degrees of separation from the epidemic origin. **A**, The results for the set D_1 before (red) and after (blue) adjusting for the hub effect. **B**, The results for the set D_2 before (red) and after (blue) adjusting for the hub effect and continuous seeding. **C**, The results for the sets D_3 and D_4 after adjusting for the hub effect and continuous seeding, and performing path reduction.

Source: Reprinted figure from Ref. [890] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

Although a single seeding event seeds the origin population with the disease, all other populations in the shortest-path tree are continually seeded by infections exported from upstream populations (Fig. 58). Such continuous seeding has been documented in the case of Zika virus in Florida coming from the Caribbean [961] and SARS-CoV-2 in the UK coming from Europe [962].

Let D_c be the set of populations that are separated by c degrees of separation from the origin population in the shortest-path tree. Let furthermore population k in D_2 be connected to origin population i via population j along the path $\psi : i \rightarrow j \rightarrow k$. After the epidemic arrives in population j at time T_{ij}^1 , population i continues to export infections to population j before the epidemic arrives in population k at time T_{ik}^1 (Fig. 58). Based on the two-population model, each imported infection that arrives in population j at times $T_{ij}^1, T_{ij}^2, \dots$ causes new exponential spreading at the hub-adjusted

rate λ_{jk} . The overall disease prevalence, $I_j(t)$, in population j at time t is therefore the sum of disease prevalence over all exponential spreadings

$$I_j(t) = \sum_{m=1}^{\infty} \mathbf{I}\{t > T_{ij}^m\} e^{\lambda_{jk}(t-T_{ij}^m)}, \quad (101)$$

where T_{ij}^m is the m th EAT in population j and $\mathbf{I}\{\cdot\}$ is the indicator function.

Based on the two-population model, the exportation of infections from population j to population k is a non-homogeneous Poisson process with the intensity function $w_{ij}I_j(t)$, which itself is a complex stochastic process due to its dependence on the random variables $T_{ij}^1, T_{ij}^2, \dots$. This leads to the probability density function of the random variable T_{ik}^n (for $n = 1, 2, \dots$), conditional on the prevalence $I_j(t)$ (and hence $T_{ij}^1, T_{ij}^2, \dots$), in the form

$$g_n(t|w_{jk}I_j) = f_{\text{Poisson}} \left[n - 1, w_{jk} \int_0^t I_j(u) du \right] w_{jk} I_j(t), \quad (102)$$

where $f_{\text{Poisson}}(\cdot, \mu)$ is the probability mass function of a Poisson random variable with the mean μ . Consequently, the unconditional probability density function of the random variable T_{ik}^n is

$$E_{T_{ij}^1, T_{ij}^2, \dots} [g_n(t|w_{jk}I_j)], \quad (103)$$

where integration proceeds over the joint probability density function of $T_{ij}^1 = t_1, T_{ij}^2 = t_2, \dots$, which in turn is given by Eq. (98) after replacing λ_i with λ_{ij} to account for the hub effect.

Ref. [890] proceeds to demonstrate that the complex dependence on $T_{ij}^1 = t_1, T_{ij}^2 = t_2, \dots$ can be simplified with a little loss of accuracy. To this end, the following certainty equivalent approximation (CEA) is made; $T_{ij}^m \approx E[T_{ij}^m | T_{ij}^1]$ for all $m > 1$. An intuitive interpretation is that most of uncertainty in the m th EAT in population j is due to uncertainty in the first EAT in this population, where the latter uncertainty is characterised by the probability density function in Eq. (92) after inserting $n = 1$ and replacing λ_i with λ_{ij} to account for the hub effect. The prevalence $I_j(t)$ becomes

$$\begin{aligned} I_j^{\text{CEA}}(t) &= \sum_{m=1}^{\infty} \mathbf{I}\{t > E[T_{ij}^m | T_{ij}^1]\} e^{\lambda_{jk}(t-E[T_{ij}^m | T_{ij}^1])}, \\ &= \sum_{m=1}^{\infty} \mathbf{I}\{t > T_{ij}^1 + \Delta T_{ij}^m\} e^{\lambda_{jk}(t-T_{ij}^1 - \Delta T_{ij}^m)}, \end{aligned} \quad (104)$$

where from Eq. (99) it follows

$$\begin{aligned}\Delta T_{ij}^m &= E [T_{ij}^m | T_{ij}^1] - T_{ij}^1, \\ &= \frac{1}{\lambda_{ij}} \exp \left(\frac{\alpha_{ij}}{\lambda_{ij}} e^{\lambda_{ij} T_{ij}^1} \right) \sum_{q=1}^{m-1} \mathbf{E}_q \left(\frac{\alpha_{ij}}{\lambda_{ij}} e^{\lambda_{ij} T_{ij}^1} \right).\end{aligned}\quad (105)$$

Finally, the unconditional probability density function of T_{ik}^n is given by

$$E_{T_{ij}^1} [g_n (t | w_{jk} I_j^{\text{CEA}})]. \quad (106)$$

Although the last expression is perfectly suitable to handle EATs for all populations in the set D_2 , additional approximations are necessary to handle populations in sets D_c , $c \geq 3$. One such approximation is *path reduction* [890] by which the path $\psi : i \rightarrow j \rightarrow k$ is treated as the direct path $\psi' : i \rightarrow k$. This allows us, for $n = 1$, to replace the probability density function in Eq. (106) with the probability density function in Eq. (92), but with suitably corrected parameters $f_1(t | \lambda_\psi, \alpha_\psi)$. The corrected parameters λ_ψ and α_ψ are obtained by minimising the Kullback-Leibler divergence (i.e., the relative entropy) [963, 890] for the first EAT through the path ψ

$$D_{\text{KL}} = \int_0^\infty E_{T_{ij}^1} [g_1 (t | w_{jk} I_j^{\text{CEA}})] \ln \frac{E_{T_{ij}^1} [g_1 (t | w_{jk} I_j^{\text{CEA}})]}{f_1(t | \lambda_\psi, \alpha_\psi)}. \quad (107)$$

The quantity D_{KL} can be understood as a measure of how much one probability distribution differs from another, reference probability distribution. By minimising the quantity D_{KL} , we therefore reduce the two leg path ψ to the one-leg path ψ' such that the probability distribution of the first EAT in population k remains unaffected. Epidemic spread from the origin population i to any population k in D_2 is thus regarded as a two-population problem, but with the local epidemic growth rate λ_ψ and the adjusted mobility rate α_ψ . If we now have an even longer path $\varphi : i \rightarrow j \rightarrow k \rightarrow m$ (i.e., $\varphi \in D_3$), we first apply path reduction to the two-leg part $i \rightarrow j \rightarrow k$, and then treat the remainder with the methods developed for the set D_2 (Fig. 59B, C).

Modelling the whole worldwide air-transportation network. To find EATs for the whole worldwide air-transportation network, it is necessary to account for the fact that each downstream population k can be connected to origin population i via multiple paths (Fig. 60A). Furthermore, paths may include

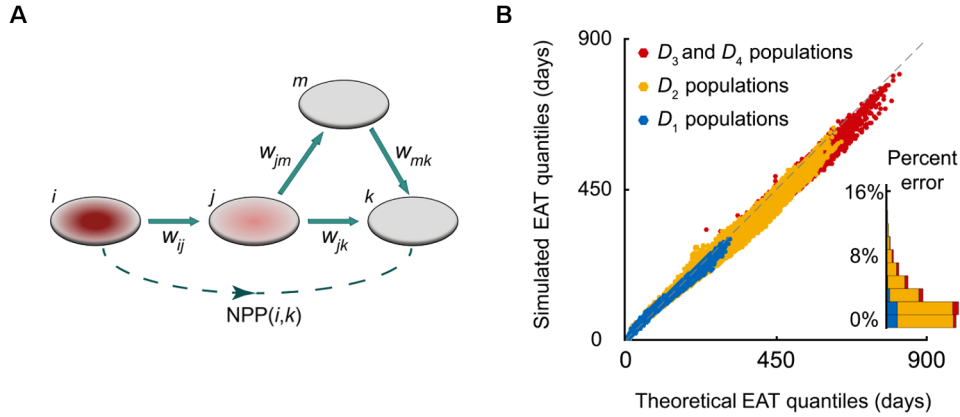


Figure 60: Numerical validation of analytics for the whole worldwide air-transportation network. **A**, Multiple acyclic paths may connect downstream population k to origin population i . **B**, Q-Q plots comparing the analytical and simulated quantiles of various EATs for downstream populations in the whole worldwide air-transportation network. Insets show the corresponding histograms of the percent error in the expected EATs. The origin of the epidemic was again assumed to be in Hong Kong. The same 100 epidemic scenarios as in Fig. 56 were used. Data points are coloured in blue for 255 D_1 populations, yellow for 1839 D_2 populations, and red for 207 D_3 and 7 D_4 populations. Source: Reprinted figure from Ref. [890] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

cycles and may intersect one another, implying that there is some degree of dependence between them. It is intuitive to assume that cycles introduce considerable delays, which makes path with cycles largely irrelevant relative to acyclic paths. If it also holds that dependence between acyclic paths is sufficiently weak to treat them as almost independent, then the following calculation becomes plausible. First, all paths that connect origin population i to downstream population k should be decomposed into a set Ψ_{ik} of ‘independent’ acyclic paths. Then, all these pseudo-independent paths are fully reduced until they are characterised by the parameters λ_ψ and α_ψ . Finally, EATs for population k can be approximated by the superposition of non-homogeneous Poisson processes [960] such that the intensity function of the superpositioned process is $\sum_{\psi \in \Psi_{ik}} \alpha_\psi \exp(\lambda_\psi t)$. Numerical results show that the entire analytical framework combining the two-population analytics, adjustment for the hub effect, adjustment for continuous seeding, path reduction, and path superposition accurately estimates EATs for almost all populations in the worldwide air-transportation network (Fig. 60B).

10.4. Bayesian inference

Bayesian inference is a class of statistical methods for data analysis and parameter estimation based on Bayes' theorem [964]. Let $P(A)$ and $P(B)$ be the probabilities of observing events A and B , respectively. Let $P(A|B)$ be the conditional probability of observing event A given the observation of event B , and $P(B|A)$ the conditional probability of observing event B given the observation of event A . Bayes' theorem says that these two conditional probabilities are related by

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}. \quad (108)$$

An analogous relationship links data to model parameters. With D denoting the observed data and θ denoting the model parameters of the data generating process, Eq. (108) can be rewritten as

$$P(\theta|D) = \frac{P(D|\theta)P(\theta)}{\int P(D|\theta)P(\theta)d\theta}. \quad (109)$$

The prior probability $P(\theta)$ represents initial beliefs about the model parameters before any data analysis. The likelihood function $P(D|\theta)$ represents the conditional probability of observing the data D given the parameters θ . The posterior probability $P(\theta|D)$ summarises the updated knowledge about the parameters upon synthesising the prior knowledge with the observed data.

The denominator of Eq. (109) requires integration over all model parameters θ , which is often very complex and analytically intractable. Numerical integration may also become computationally prohibitive as the number of parameters increases. However, the denominator only plays the role of a normalisation constant. Sufficient for inference is the proportionality

$$P(\theta|D) \propto P(D|\theta)P(\theta). \quad (110)$$

Accordingly, Bayesian inference mainly comprises two steps; (i) formalising the prior distribution of each model parameter using background knowledge and literature reviews and (ii) designing the likelihood function by using probabilistic models to account for the underlying data-generating process.

Specifying prior distributions is a nontrivial task [964]. Existing studies in the fields of infectious diseases modelling, network theory, bioinformatics, and statistical physics tend to use the simplest uninformative flat or diffuse

prior [965]. The main aim of such a simplification is to assess the capacity of the likelihood model in fitting the observed data. Despite being useful in resolving low-dimensional problems with a few parameters to infer, the flat or diffuse prior cannot be regarded as a universal tool for fitting. In particular, for high-dimensional problems with many parameters to infer, the usage of flat or diffuse priors can lead to biased estimations or convergence failures [964]. Recent progress in the field suggests that even weakly informative priors may be a better option. How to setup prior distributions is explained in detail in Refs. [966, 967].

To estimate model parameters, Bayesian methods often formulate the likelihood function using probabilistic models. The purpose of such models is to describe the data-generating process behind observed data. In this section, we proceed by outlining two epidemiological case studies to explain how to develop the likelihood function using probabilistic models for Bayesian inference.

Inferring the basic reproductive number R_0 for the 2009 influenza A(H1N1) in Greater Mexico City. As briefly discussed at the beginning of this section, the basic reproductive number R_0 is an important epidemiological quantity, giving the expected number of secondary infections induced by an infectious person in a fully susceptible population. Estimating R_0 during the early stage of an outbreak is key to understanding the potential of the disease for interpersonal transmissions, the requirements for a vaccine to achieve herd immunity, and the extent of non-pharmaceutical interventions to control the outbreak.

Here, we look at the estimation of R_0 in the case of the 2009 influenza A (H1N1) epidemic in Greater Mexico City. Ref. [968] is well-known for presenting a maximum likelihood method that was used to estimate R_0 in this particular case. The method runs a large number of computer simulations to explore the parameter space, which is computationally so intensive that it requires the use of a supercomputer. Ref. [890] lessens the computational burden by using Bayesian inference in conjunction with disease-exportation records from Mexico to the first 12 countries as summarised in Ref. [968]. This inference combines the two-population model in Eq. (92) with adjustments for the hub effect in Eq. (100).

The estimates of the basic reproductive number R_0 for the 2009 influenza A(H1N1) using the GLEAM simulator powered by high-performance computing [968] equal 1.65 with the 95%-confidence interval (CI) [1.54, 1.77],

1.75 with the 95 % CI [1.64, 1.88], or 1.89 with the 95 % CI [1.77, 2.01] depending on whether the outbreak in La Gloria, Mexico, started on 11, 18, or 25 February 2009, respectively. Bayesian inference, by contrasts, rests on a likelihood function that can be written in a closed form that only depends on the basic reproductive number R_0

$$L(R_0) = \prod_{j \in A} f_1(t_j | \lambda_{ij}, \alpha_{ij}) \prod_{j \in B} F_1(t_j | \lambda_{ij}, \alpha_{ij}), \quad (111)$$

where population i denotes the Greater Mexico City as the epidemic origin, t_j denotes the observed EAT for population j which can be exact (set A) or left-censored (set B), $\lambda_{ij} = \lambda_i - \sum_{k \neq j} w_{ik}$ denotes the hub-adjusted epidemic growth rate, and α_{ij} denotes the adjusted mobility rate. It holds that $R_0 = 1 + \lambda_i Z$, where Z is the mean infectious period. This likelihood function leads to the same estimates of R_0 as the GLEAM simulator but without relying on high-performance computing (Fig. 61). Such a substantial reduction in computational complexity and resource requirements is expected to greatly improve the efficiency and timeliness of pandemic forecasting in the future.

Estimating the transportation risk of Covid-19 from Wuhan to other cities in China. Covid-19 is caused by the SARS-CoV-2 virus. Due to the rapid global expansion of this virus, rising death numbers, unknown animal reservoir, and the increasing evidence of interpersonal transmissions [969], the World Health Organization (WHO) declared a public-health emergency of international concern on 30 January 2020. An important concern at the time was the risk of new cases spreading from Wuhan to other locations. This concern was addressed in Ref. [955], a quick case study performed during early 2020 using the methods proposed in Ref. [890].

Estimating the risk of new cases spreading from an origin population outwards begins with an epidemiological model. Let $\Delta I_W(t)$ be the daily number of new infections in Wuhan from 1 December 2019 through 22 January 2020. Based on the epidemiological data from the first 425 Covid-19 cases confirmed in Wuhan by 22 January 2020 [969], the epidemic was assumed to grow exponentially

$$\Delta I_W(t) = i_0 \exp(\lambda t), \quad (112)$$

where i_0 is the number of initial cases on 1 December 2019, and λ is the local epidemic growth rate between 1 December 2019 and 22 January 2020. New Covid-19 cases were typically detected with a mean delay of $D = 10$

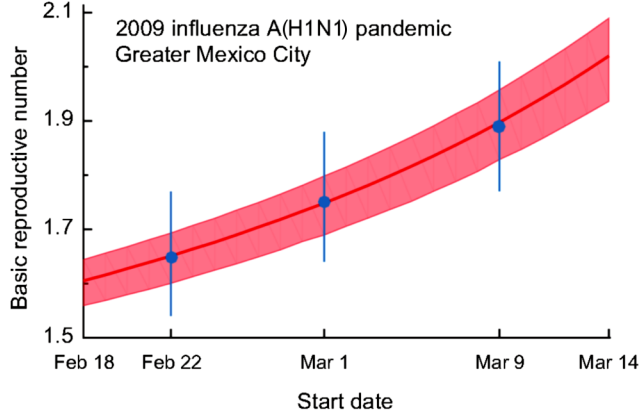


Figure 61: Inference of the basic reproductive number R_0 for the 2009 influenza A(H1N1) pandemic in Greater Mexico City. The value of R_0 is inferred using EATs for the first 12 countries seeded by Mexico, as documented in Ref. [968]. The red curve and the red-shaded area respectively indicate the posterior medians and the 95 % credible intervals. The blue dots and error bars respectively show the mean and the 95 % confidence intervals from Ref. [968] depending on whether the influenza A(H1N1) pandemic started in La Gloria, Mexico on 11, 18, or 25 February 2009.

Source: Reprinted figure from Ref. [890] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

days [970], which included an incubation period of 5-6 days [969] and a delay from symptom onset to detection of 4-5 days. With this in mind, the number of infectious cases at time t is given by

$$I_W(t) = \sum_{u=t-D}^t \Delta I_W(u), \quad (113)$$

and the prevalence of infectious cases is

$$\eta(t) = I_W(t)/N_W, \quad (114)$$

with $N_W = 11.1$ million denoting the population size of Wuhan.

The next step in estimating the risk of new cases spreading from an origin population outwards entails specifying a model of mobility. Assuming that the visitors to Wuhan and the residents of Wuhan share the same daily risk of infection, a non-homogenous Poisson process can be used to estimate the risk for exporting Covid-19 infections from Wuhan [957, 958, 890]. Let $W_j(t)$ be the number of Wuhan residents travelling to city j on day t , and $M_j(t)$

the number of travellers from city j travelling back from Wuhan on the same day. The intensity function of the non-homogeneous Poisson process is then $\eta(t) [W_j(t) + M_j(t)]$, and the probability of introducing at least one Covid-19 case from Wuhan to city j by time t is

$$1 - \exp \left(- \int_{t_0}^t \eta(u) [W_j(u) + M_j(u)] du \right), \quad (115)$$

where t_0 is the start of the study period, that is, 1 December 2019.

To estimate the unknown parameters, such as the number of initial cases i_0 and the local epidemic growth rate λ , a likelihood function is needed. This function can incorporate diverse information, including the information on disease exportations outside of China even if risk is to be quantified solely for Chinese cities. Ref. [955] in particular used the data on EATs due to 19 Wuhan residents who travelled to 11 cities outside of China before 22 January 2020.

Let N_j be the number of infectious Wuhan residents detected at location j outside of China, and t_j^i the time at which the i th detection occurs. Furthermore, let t_j^0 denote 1 January 2020, which is the date on which international surveillance for infected travellers from Wuhan began. Finally, let t_e denote 22 January 2020 which is the end of the study period. As mentioned above, the rate at which infected residents of Wuhan arrive at location j at time t is $\eta(t)W_j(t)$. Accordingly, the likelihood function for observing the EATs due to 19 Wuhan residents travelling outside of China by 22 January 2020 is

$$\prod_{j=0}^{11} \exp \left(- \int_{t_j^{N_j}}^{t_e} \eta(t)W_j(t)dt \right) \prod_{i=1}^{N_j} \eta(t_j^i)W_j(t_j^i) \exp \left(- \int_{t_j^{i-1}}^{t_j^i} \eta(t)W_j(t)dt \right). \quad (116)$$

The first product quantifies the probability of not seeing any exportations at location j between the time of the N_j th exportation and the study end. The second product quantifies the probability density of seeing the i th exportation at location j at time t_j^i . All cities included in the study but without observed Covid-19 cases before 22 January 2020 were treated as a single location indexed by $j = 0$.

With the likelihood function in Eq. (116), it becomes possible to estimate the number of initial cases i_0 on 1 December 2019 and the local epidemic

growth rate λ between 1 December 2019 and 22 January 2020. Ref. [955] used the Markov Chain Monte Carlo method with the Hamiltonian Monte Carlo sampling and non-informative flat priors. By assuming that the incubation period is exponentially distributed with the mean L of 3-6 days and the infectious period is also exponentially distributed with the mean Z of 2-7 days, the basic reproduction number is $R_0 = (1 + \lambda L)(1 + \lambda Z)$. With all the parameters estimated, Eq. (115) calculates the risk of transporting at least one case from Wuhan to a downstream city j before the lockdown of Wuhan on 23 January 2020 (Fig. 62).

Ref. [955] estimates the Covid-19 doubling time ($= \ln(2)/\lambda$) at 7.31 days with the 95 % credible interval (CrI) [6.26, 9.66] days. Other studies using similar methods have yielded congruent results; for example, Ref. [956] estimates the doubling time at 6.4 days with the 95 % CrI [5.8, 7.1] days and the basic reproductive number R_0 at 2.7 with the 95 % CrI [2.5, 2.9]. Ref. [971] presents a metapopulation network model covering 375 Chinese cities and employs Tencent migration data to capture population movements during the 2021 Spring Festival period in China. By fitting the model to the reported 801 Covid-19 cases throughout China after the lockdown of Wuhan on 23 January 2020, the basic reproductive number R_0 is estimated at 2.38 with the 95 % CrI [2.03, 2.77]. Interestingly, Ref. [972] estimates the basic reproduction number during the first wave of Covid-19 in mainland China at $R_0 = 5.7$ with the 95 % confidence interval [3.8, 8.9], which is inconsistent with other studies. This extreme result, however, may be due to an improper assumption of a single infection occurring at the initial time.

10.5. Challenges and future work

We have taken a look at the state of the art in epidemiological modelling and how it relates to the budding field of digital epidemiology. Here, we outline some pressing issues and ideas for further progress.

The metapopulation approach described herein considers mixing in each of the populations (i.e., cities) to be homogeneous. However, looking at mobility patterns only on the intercity scale hides away important epidemiological phenomena that happen on the intracity scale [973, 974]. Ref. [975], for instance, describes a metapopulation susceptible-exposed-infectious-removed (SEIR) model that incorporates intracity mobility to explore the spread dynamic of Covid-19 in ten metropolitan areas in the US. Doing so has enabled identifying higher infection rates among disadvantaged racial and socioeconomic groups because of the differences in mobility. Specifically, disadvan-




Figure 62: Risks of exporting Covid-19 from Wuhan, China, before the lockdown on 23 January 2020. **A**, Daily travel volume to and from Wuhan. **B**, Estimated and confirmed cumulative Covid-19 cases in Wuhan. Green line and grey shaded area indicate the mean and the 95% credible interval (CrI) of the estimated cumulative true infections since 1 December 2019. Black dots indicate the cumulative confirmed-case counts during 1-22 January 2020. January 10 marks the beginning of the Spring Festival travel season in China. **C**, Probabilities that Chinese cities import more than one Covid-19 case from Wuhan by 22 January 2020. 131 cities (orange and red circles) exhibited a high risk of more than 50%.

Source: Reprinted figure from Ref. [955] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

taged groups have relatively little control over reducing their mobility and consequent exposure to infectious diseases.

In addition to mobility patterns, the epidemiological importance of contact patterns is impossible to overlook. These latter patterns have been shown to be highly assortative with age; especially school children and young adults tend to mix with similarly aged people [976, 977, 978]. This leaves younger populations potentially more vulnerable to infectious diseases unless there are attenuating biological circumstances, as is the case with (the early variants of) SARS-CoV-2 [979].

Another key factor in epidemiology is individual heterogeneity in the ability to transmit an infectious disease. Superspreaders, for example, cause disproportionate number of secondary cases during the outbreaks of measles, influenza, rubella, smallpox, Ebola, monkeypox, SARS, and Covid-19. In the case of Covid-19, about 19% of infectious individuals seed 80% of all local transmissions [980]. SARS-CoV and SARS-CoV-2 viruses have, in fact, recently been shown to cause the number of secondary infections that follows a fat-tailed distribution, thus emphasising a large heterogeneity in transmission ability among individuals [981]. Accounting for this heterogeneity in epidemiological models generates the results that differ substantially from average-based approaches such that outbreaks are rarer but more explosive [982]. Future models should therefore account not only for heterogeneity in contact networks, but also heterogeneity in infectiousness and possibly other epidemiological parameters.

Since the onset of the Covid-19 pandemic in December 2019, the amount of research on non-pharmaceutical interventions has exploded [983]. A big reason for such an explosion of interest is that non-pharmaceutical interventions are the only means of reducing the spread of a novel pathogen. They are also effective. For example, social distancing alone has proven sufficient to control Covid-19 in China [984]. Fast isolation of infectious individuals has furthermore reduced the time between successive disease onsets in a transmission chain (i.e., the serial interval), thus signalling effectiveness in preventing multiple secondary infections that would have arisen without isolation [985]. These examples show that determining the optimal combinations of non-pharmaceutical interventions in given circumstances may save many lives, and should therefore be a major component of epidemiological modelling.

This chapter has hopefully demonstrated just how much data-hungry epidemiology is. A major limitation in the use of digital-data sources, such as

Google Trends, Twitter, and Facebook, is that that none of them have been constructed with epidemiology in mind. This limitation can be overcome by establishing data standards and specialised systems akin to HealthMap [942] and Influenzanet [986], but also proving the usefulness of these systems to public-health agencies. Achieving so is no small feat because complicated data-analysis and modelling methodologies are often too demanding for otherwise busy public-health officials. In this context, relying on easy-to-use interactive interfaces called visual analytics [987] may help. Ref. [988], for example, demonstrates how to effectively couple data mining and agent-based epidemic modelling with a visual-analytics environment to facilitate human decision making in controlling infectious diseases.

11. Environment

Anthropogenic impact on the geological record is such that the 20th century saw a start of a new geological epoch—Anthropocene [989]. While climate change may be the (politically) most prominent issue today, our civilisation has additional profound impacts on every single ecosystem through pollution and habitat loss. Surprisingly, these are ultimately social issues because only through societal consensus on the need for action can they be controlled [990].

Building social consensus against environmental degradation is extremely difficult. Preventing environmental degradation incurs unwanted costs either through the need for direct investment into unprofitable infrastructure and processes (e.g., water treatment facilities, carbon capture, safe waste disposal), or through opportunity costs (mostly land use restrictions due to ecosystem conservation, e.g., nature protected areas). Hence, to stay competitive, economies generally allow externalisation of environmental degradation costs until it is proven that costs of environmental degradation outweigh their economic benefits. Even once proven, building the social consensus can be extremely difficult; climate change and plastic pollution are just the two most prominent current cases that highlight the difficulties.

Despite the difficulties, change is possible if sufficiently strong scientific case can be made, and costs of degradation can be quantified. For example, theory on the adverse effect of at least some chlorofluorocarbons (CFCs) on the protective ozone layer was settled in 1974 [991], and experimentally confirmed in 1985 [992]. It took only two years following the indisputable evidence for international ratification of a treaty to phase out the use of

ozone-depleting substances [993], and the ozone layer is recovering [994]. Tetraethyllead in gasoline followed a similar path from suspicion in mid-1920s [995] to ban in 1990s and 2000s as scientific evidence pioneered by Patterson [996] accumulated. Due to largely local nature of the problem, nations were able to take regulatory steps independently, with Japan banning low-octane lead gasoline in as early as 1975 [997], and some countries continuing the use to date. Similar stories can be told of numerous other chemicals ranging from mercury, to dichlorodiphenyltrichloroethane (DDT), to polychlorinated biphenyls (PCBs) and polybrominated diphenyl ethers (PBDEs).

History—and the present—clearly demonstrate that the burden of proof that an externality is overly damaging squarely lies with environmental science, declaratory proclamations of the precautionary principle notwithstanding. Furthermore, the weight of the evidence required to elicit science-based activism that may eventually lead to societal consensus and (ultimately) solutions is extremely high. To provide sufficiently strong evidence, environmental science had to evolve, mostly towards physics. Because climate change will be—due to its importance—discussed in a separate section (Section 12), this section will focus on pollution and physical habitat loss. Historically, physical habitat loss due to competition of wildlife for natural resources with humans has been the chief driver of extinctions; pollution as an environmental problem, however, is a relatively new phenomenon. Before the industrial age, humans simply could not (and had no motivation to) produce toxic chemicals at a scale that could severely impact wildlife.

11.1. Pollution

There are many types of pollution, all traceable to introduction of materials or energy into the environment. The materials range from simple to complex chemicals, to particles, or even displaced natural materials. Energy pollution can also be diverse, ranging from noise, to non-ionising electromagnetic radiation (including light), to ionising radiation. Judging impact of pollution can be challenging because the judgement depends on both the point of view, and on the existing knowledge. For example radioactive debris from the Chernobyl nuclear plant explosion were not negative from the point of view of wildlife; in fact, wildlife is thriving despite the moderately toxic environment [998] simply because anthropogenic influence before evacuation was even worse. Focusing on the organismal level helps minimise such ambiguities.

Traditionally, toxicity has been estimated by exposing model organisms to varying concentrations and doses of chemicals in food or the environment while tracking endpoints like survivorship or mortality, fertility, cognitive ability, and—more recently—biomolecular targets like gene expression, protein levels, and reactive oxygen species. Such data yields dose-response curves used to predict the no-effect concentration, that is, the environmental concentration at which no negative effects are expected. Some calculations of the concentration simply divide by 1000 the concentration at which 50 % of organisms in short-term experiments showed a response, such as death or some other endpoint [999].

Considering the consequences of getting it wrong, the inadequacy of testing and observations as bases for regulation is staggering:

- Species or even individuals have different sensitivities to toxic exposure; the precautionary principle demands that generally the most sensitive model organisms should be used [999], but the choice of test organism is limited because in general only animals thriving in laboratory environments can realistically be utilised, and endangered species cannot be used at all.
- Exposures in the lab are standardised, and typically last only a small fraction of organisms' life span [1000]. Consequently, multigenerational studies are extremely rare and effects of long-term or cross-generational exposure are rarely captured.
- Environmental conditions such as food availability, temperature, and humidity change exposure and its effects [1001]. Relevance of tests in standardised laboratory environments to effects in highly variable natural environments is, therefore, limited.
- Ecological feedbacks that concentrate toxicants, such as biotransformation and bioaccumulation, can result in much higher exposure for some species than indicated by experimental observations [1002].
- Even though chemicals mix in the environment, and mixtures can exacerbate toxicity [1003], most chemicals are tested—and their legal limits set—independently.
- Only a small fraction of new chemical compounds are tested for toxicity. More than 20 million substances were reported as of January 2017, with

about a million compounds added annually at an exponentially growing rate [1004]; number of tested chemicals is measured in thousands.

Accordingly, relying on testing for legislation is highly impractical at best, misleading or wrong at worst. Consequences of underrating dangers may be dire—as previous examples show, it is extremely difficult to eliminate impacts of a chemical once it permeates the environment. Microplastics [1005] and whatever is causing the decline of pollinators [1006] are shaping up to be the next major problems, but each new compound presents a new risk by itself or in synergy with other pollutants. High-throughput screening could help identify the most dangerous chemicals and direct preventive research in the future [1007], but process-based modelling is leading the way in predictive ecotoxicology.

Process-based models rely on physical principles to explain observations, predict effects of hitherto unobserved exposures, and capture interactions between organisms and the environment. Although process-based models cover the whole range of scales, from molecular to landscape and ecosystem, the scales are not sufficiently linked [1008].

For example, quantitative structure-activity relationship (QSAR) models link characteristics of chemicals to their physicochemical, biological, and environmental properties. While essentially correlative in nature, QSARs use mechanistic descriptors derived from physico-chemical properties of chemicals. Types of the properties used to derive the descriptors determine QSAR dimension: chemical formula (0D), sub-structural fragment (1D), graph theory (2D), spatial geometry (3D), conformation/orientation/protonation state (4D), induced fit on ligand-based virtual or pseudo-receptor model (5D), 5D plus other solvation conditions (6D), and real target receptor model data (7D). In particular, descriptors in 3D QSARs involve a number of concepts from physics such as energy minimisation, classical approaches to molecular mechanics, quantum mechanics implemented using Born-Oppenheimer or Hartree-Fock approximations, or density functional theory to reduce computational loads and reduce descriptors to quantitative variables suitable for the correlative step in QSAR analysis. The correlative step, that is, relating descriptors to outcomes, can utilise any of the number of statistical methods ranging from multiple regression to artificial neural networks.

Despite their proven track record, QSARs have limitations. Due to the statistical nature of development and rigorous validation requirements, QSARs are still very data-intensive. Furthermore, when used to predict

high-level responses (such as mortality), QSARs offer limited insight into metabolic pathways of toxic action. Understanding these, however, can be crucial for considering effects of chemicals in untested systems and mixtures. To overcome these shortcomings, the Organisation for Economic Co-operation and Development (OECD) actively supports development of a modular framework of metabolic cause-effect transfer functions: adverse-outcome-pathway (AOP) framework [1009].

The AOP framework captures effects of exposure by modelling a sequence of all relevant molecular and cellular events. The framework consists (as of January 2021) of more than 2,000 *Key Events* (KEs). Exposure to a stressor triggers a KE that assumes the status of a *Molecular Initiating Event* (MIE), which initiates a chain of KEs through a series of *KE Relationships* (KER) akin to if→then prescriptions. The chain terminates with an *Adverse Outcome* (AO). For example, in AOP #15 any one of 10 stressors recognised by the framework can cause DNA alkylation (MIE), which through KER #24 leads to inadequate DNA repair (KE #155). KE #155 then through KER #164 leads to KE #185 (increase in mutations). The increase, through KER #202, leads to a heritable increase of mutations in offspring (AO #336). Elements of AOP #15 relate to seven other AOPs. Each element in the AOP chain has detailed documentation, a rigorous scientific background, and has been reviewed by experts.

Due to stringent background and review requirements, only 16 of more than 300 existing pathways have been endorsed across the OECD, but the framework is expected to expand exponentially as new pathways reuse old events and relationships, thus requiring fewer new ones. Currently, the framework focuses on humans, but similarity of organisms on bio-molecular level offers hope that eventually other organisms could benefit as well. High-throughput screening could provide a fast and affordable way to determine new key events and relationships. Currently, AOPs are not suitable for environmental ecotoxicology, making Dynamic Energy Budget (DEB) model the tool of choice for advanced links between environmental exposure and ecologically relevant organism-level endpoints.

DEB theory [8] is, essentially, an application of the laws of thermodynamics to all three types of macrochemical reactions of a heterotrophic aerobe: assimilation, growth, and dissipation [9]. Focusing on (i) the four building blocks constituting 99% of living biomass (C, H, O, and N), and (ii) effects of a fairly limited number of ‘hub’ metabolites crucial for metabolic-network function that are markedly similar between species on macromolecular and

cellular levels, the theory makes a number of simplifications leading to a robust theoretical framework able to capture ontogeny of living organisms and make testable predictions. Standardised DEB models are aggregated in the AmP database [1010], which at the moment of writing contains entries for over 2,800 species. Detailed derivation of the standard DEB for the physics-minded reader is given in Ref. [9].

Starting from the first law of thermodynamics applied to an organism, the rate of change in internal energy U is:

$$\frac{dU}{dt} = \dot{Q} + \dot{W} + \sum_i \bar{h}_i \left. \frac{dM_i}{dt} \right|_{\text{in}} - \sum_i \bar{h}_i \left. \frac{dM_i}{dt} \right|_{\text{out}}, \quad (117)$$

where \dot{Q} and \dot{W} respectively represent heat-transfer rate and mechanical power, $i \in \{X, P\}$ stands for organic substances in food and faeces, $i \in \{C, H, O, N\}$ stands for metabolites, \bar{h}_i are molar enthalpies, and M_i is the amount of substance i in the organism. Next, using the mass, energy, and entropy balances, it can be shown that the total organismal Gibbs free energy is a sum of Gibbs free energies in compartments of the organism, assuming homeostasis of each compartment, that is, the chemical composition of each compartment remains constant throughout the life of the organism. The standard DEB theory recognises two compartments: energy reserve and structure, with an additional compartment tracking energy committed to maturation and reproduction (Fig. 63). If the organism is isomorphic (i.e., of constant shape), energy reserve and structure in a standard DEB model can be described by a simple set of coupled ordinary differential equations, especially when scaled

$$\frac{de}{d\tau} = g \frac{f - e}{l}, \quad (118a)$$

$$\frac{dl}{d\tau} = \frac{g}{3} \frac{e - l - l_T}{e + g}, \quad (118b)$$

where $0 \leq f \leq 1$ is the scaled functional response representing surface-specific assimilation rate relative to the maximum, e is the energy density of reserve relative to the maximum, l is the length of the organism relative to the maximum possible length for $f = 1$, g is a compound parameter called the energy investment ratio, and l_T is the heating length, scaled parameter accounting for energy spent on maintaining target body temperature in endotherms. In constant environments, the standard DEB model converges to




Figure 63: Basic metabolic processes of heterotrophic aerobes according to standard DEB theory. Food is assimilated into reserve. In turn, the reserve is (i) converted into structure representing growth, (ii) committed to reproduction when possible, and (iii) used to power various dissipative processes such as maturation, maintenance, and metabolic inefficiencies (i.e., overheads) of growth, assimilation, and reproduction. Non-assimilated food (faeces) and other excess metabolites such as carbon dioxide, water, and nitrogenous waste are excreted into the environment.

the von Bertalanffy growth equation [1011], which is still the most widely used energy-based equation for organismal growth. The von Bertalanffy growth equation is, however, a demand-side model; it can help estimate energy required for observed growth, but cannot predict how growth (or any other processes) would respond to changes in the environment.

DEB models, on the other hand, make a quantum leap in that they causally link environmental energy and material availability to organismal growth and reproduction whilst preserving mass and energy balances. Additionally, the standard DEB theory also enables tracking of metabolism- and stress-related hazard rate, that is, the risk of death due to accumulated (cellular) damage. These features make some extraordinary feats possible. For example, the ability to track material fluxes in the context of interactions between the environment and the organism enabled a revolution in toxicokinetic (the distribution of toxicants) and toxicodynamic (toxicant effects) modelling by (i) resolving a number of long-standing issues of empirical dose-response curves [1012], and (ii) enabling prediction of bacterial population dynamics under exposure seven times greater than used for model fitting [1013] (Fig. 64), as well as identification of nano-toxicity and its mechanisms [1014].

Further exemplary successes achieved by mechanistic modelling advocated by DEB theory include reconstructing otolith growth in anchovy [1015, 1016] that enables tracking of historic environmental conditions and therefore determination of organism's population range (Fig. 65). Ecological interactions can explain why a level of exposure that measurably harms a plant




Figure 64: Cadmium-ion toxicity. DEB model for Cd-ion toxicity predicts bacterial population dynamics for exposures of up to 150 mg(Cd)/L (solid curves) with a single common parameter set fitted only using observations at 0, 10, and 20 mg(Cd)/L (dotted curves). Source: Reprinted figure from Ref. [1013] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

can increase its growth and yield (Fig. 66). Finally, the concept of hazard rate unifies the Weibull (allometric) and Gompertz (exponential) models of ageing, and is able to explain why hungry mice live longer (Fig. 67).

DEB makes strides in understanding toxic effects of exposure on individual and population levels, but falls short of the ultimate goal of ecotoxicology: understanding multi-generational effects of toxicants in environmental settings where multiple toxicants combine and interact with other stressors, and where complex ecological interactions could greatly affect outcomes. Reaching the goal requires modelling multiple populations exposed to a variety of toxicants in heterogeneous (spatially explicit) environments.

Many building blocks necessary to reach said ultimate goal of ecotoxicology already exist, and tools reduce barriers to entry. A conceptual framework has been developed that links the sub-cellular and sub-organismal AOP framework to individual-level DEB models [1008]. *General Unified Threshold model for Survival* (GUTS) is a modelling framework for toxicity test analysis in which ‘survival’ is the endpoint [1018], with open source software toolbox available at openguts.info. The framework has been recognised in the Or-




Figure 65: Comparison of real and simulated otoliths. **A**, Seasonal variability in opacity patterns. Observed (dashed) and simulated (solid) variability in opacity of southern North Sea (NS, black) and Barents Sea (BS, red) anchovy otoliths. **B**, Opacity images: actual (top left) and simulated (bottom left) North Sea anchovy otoliths, and actual (top right) and simulated (bottom right) Barents Sea anchovy otoliths. Of note is that only environmental forcing (temperature and food) differ between the two populations; the model and parameter values are equal in both simulations. Source: Reprinted figure from Ref. [1015] under the Creative Commons Attribution 4.0 International (CC BY 4.0).

ganisation for Economic Co-operation and Development (OECD) guidance for toxicity testing since 2006 [1019]. Effects of toxicant mixtures have successfully been estimated from effects of each toxicant alone [1020, 1021, 1022].

DEBKiss is a simplified version of the standard DEB model that, at the expense of generality, substantially reduces the barrier to entry to DEB modelling whilst preserving much of the utility, especially for specific questions for which cross-species comparison is not of primary importance [1023]. Older simple energy budget models exist, most notably net-assimilation and net-production models [1024]. These too have been proven useful in specific

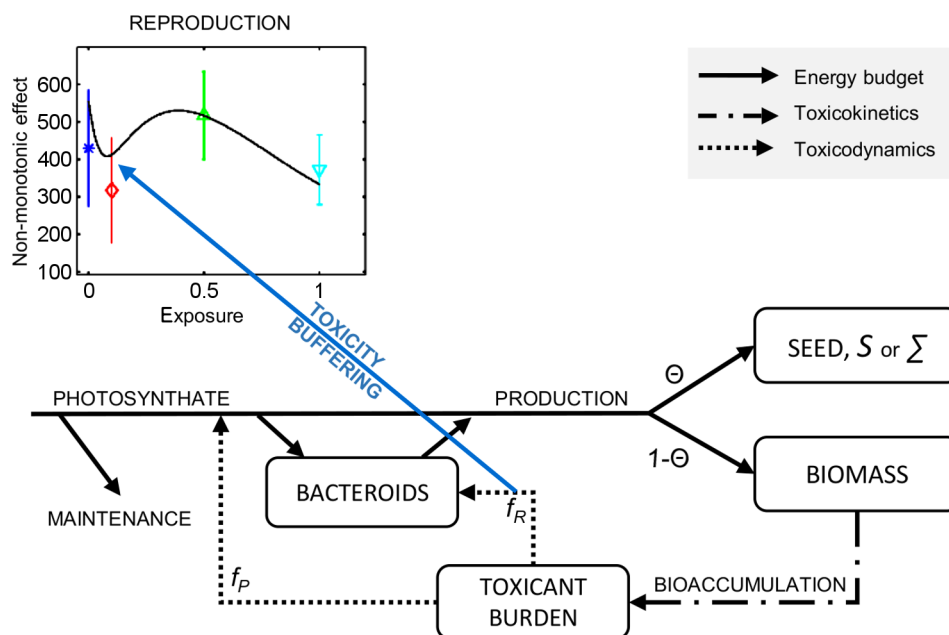


Figure 66: Ecological interactions affect the outcome of pollution. Despite clear soybean plant cellular damage caused by exposure to CeO_2 nanoparticles and negative effects of low exposure to growth and yield, higher exposures surprisingly improve plant growth and yield (inset). DEB model of coupled plant-bacteroid dynamics provides an explanation. Photosynthate (energy) is utilised for (i) maintenance of the plant, (ii) growth of bacteroids; remaining energy is used for (iii) seed production with proportion Θ , and plant growth with proportion $1 - \Theta$. Bioaccumulated toxicant affects both the bacteroids (f_R) and the plant (f_P). In nitrogen-poor soil, the bacteroids provide nitrogen needed for plant growth, but in nitrogen-rich soil, the bacteroids reduce energy available to the plant without providing any benefits. Small exposure affects the plant without killing bacteroids, causing a depression in growth and yield. Higher exposures, however, kill off the bacteroids but not the plant. This leaves more energy for the plant, thus improving growth and yield. See Ref. [1017] for further details.

circumstances, but lack scientific rigour and generality of the standard DEB family of models.

DEBTool is a set of Matlab scripts to estimate parameters of a DEB model and run it. The tool is open source. Accompanying this tool is *Add-my-Pet* (AmP) database of DEB model parameters containing parameters and data for over 2800 species and growing. Ref. [1010] provides a good overview of the database functionality.

DEB models are naturally suited for *individual based modelling* (IBM) in

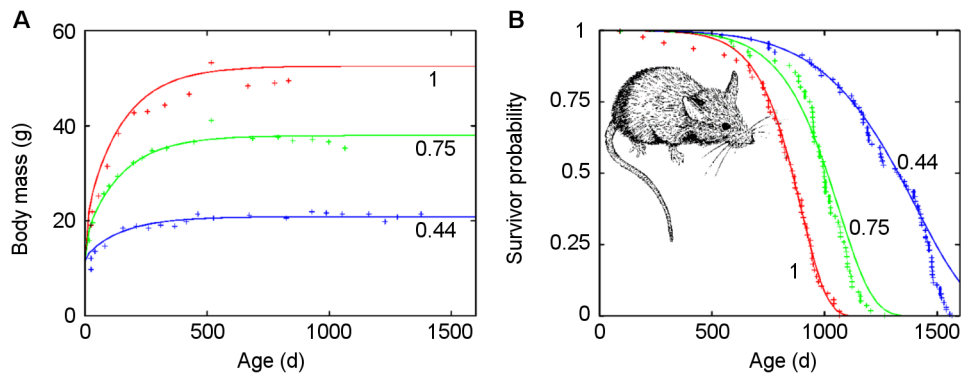


Figure 67: Longevity of mice depends on food abundance. **A**, Growth of mice. **B**, Survivor probability. Solid curves (model simulations) and data ('+') are shown for food abundances of 0.44 (blue), 0.75 (green), and 1 (red) relative to the maximum. Mice at restricted diets have a smaller metabolic activity that produces less damage-inducing compounds, thus accumulating less damage per unit of time, and living longer. See Ref. [8] for further details.

Source: Courtesy of Sebastiaan A. L. M. Kooijman.

which each organism is modelled separately; Ref. [1025] provides a Graphical User Interface (GUI) in NetLogo for simple IBM-based population modelling. The standard DEB model can be run directly in the GUI, but more complex features—including spatial heterogeneity—require additional programming. Population dynamics can be modelled in a number of ways if IBM approaches are impractical, including the Euler-Lotka equation [1026, 1027], matrix population models [1028, 1029], continuous-time physiologically structured [1030, 1031], and integral-projection models [1032]. Of those, matrix population models have a particularly low barrier to entry and can also separate the population into patches [1033], include predation, and other ecological interactions. Escalator Boxcar Train tool (EBTtool) is a GUI-based environment for implementation and analysis of physiologically structured models based on any physiological model of an individual including DEB.

Maxent [1034] is a Java application that utilises data on species occurrences and environmental conditions to predict the geographic distribution of species using a maximum entropy approach. The software, appearing in thousands of publications, has been open source since 2017 [1035]. *NicheMapR* is an R package available at [mrke.github.io/](https://github.com/mrke) that serves roughly the same purpose as Maxent, but is substantially more advanced. While Maxent integrates only observation data, NicheMapR is able to integrate metabolic

models such as DEB, account for heat and water exchange using principles of biophysical ecology including behaviour, as well as for microclimate—using micrometeorology, soil physics, and hydrology—to which the organism is exposed [11].

Therefore, frameworks and tools linking scales of biological organisation already exist, with physics providing the necessary glue. Crucially, frameworks make predictions that can be falsified, thus respecting the scientific process. Due to their mechanistic origins, the tools and frameworks are largely modular and can therefore be used to investigate effects of other anthropogenic pressures, including climate change and habitat loss.

11.2. Physical habitat loss

Physical habitat loss happens when a domestic species can no longer inhabit an area. Anthropogenic pressures, including climate change, can drive habitat loss through change of local environmental conditions that can facilitate exclusion of a domestic species by introducing—or merely making the local environment better for—an invasive species. Most habitat loss to date, however, has been due to land and sea use, that is, direct competition between wildlife and humans for space and related natural resources such as water and, in the case of predators, prey biomass. Measurable human influence on the environment may have started millions of years ago [1036]; today, 95 % of Earth’s land masses show human influence, with the remaining 5 % being mostly in inhospitable areas, including ice [1037].

Unlike human-driven extinctions, the idea that we might need to preserve at least some of the natural biodiversity is relatively new. First legal environmental protections focusing on ecosystem preservation sprung up in the second half of the 18th century, and significantly proliferated only in the 20th. The idea that environmental impact should be assessed at least for large-scale projects is even newer. The USA was in 1969 the first to require environmental impact assessments (EIA) for large-scale projects. To date, protected areas and the assessments remain the most effective tools for habitat preservation and biodiversity conservation; the number of protected areas has rapidly grown in the past decades, and the range of projects required to have an environmental impact assessment has been greatly expanded in the developed world.

To date, 15.4 % of terrestrial areas and 7.6 % of marine areas are classified as protected. These include protected areas ranging from strict nature preserves to airports [1038]. Strict nature preserves where no human activities

are allowed except scientific monitoring and management interventions are rare. The vast majority of protections allow at least some activities. Therefore, management of protected areas involves balancing human activities with nature preservation.

More often than not, such management is focused on maximising human activities while keeping indicators of environmental damage within acceptable limits. Hence, protected areas are managed to a level of human activities just shy of environmental destruction, that is, to maximum damage levels that the environment can sustain in the long term. Clearly, this is not ideal even if the managing authority has the greenest of intentions. What we may perceive as the safe level of damage based on current data may reduce the resilience of the protected ecosystem to the point at which it succumbs to new pressures, such as the climate change. Tourism in national parks is a poster child for such a balancing act. Because nature is the main product offered to tourists, managing authorities have a vested interest in preserving the environment. Nevertheless, visitations are often maximised at the expense of nature [1039].

Similar considerations hold for fisheries, with the added complication that the fishing industry has to contend with both the regulatory and natural uncertainties. For example, EU sets total allowable catch on (bi)annual basis to preserve fish stock and biodiversity. Without public subsidies, a particularly low value could decimate the fishing fleet because the industry may not be able to absorb the reduction in earnings. Such a reduction carries negative social and economic implications, but also makes the industry unable to respond to a higher allowed catch in the future. Hence, the regulatory authority has to balance the industry's need for income stability with conservation goals, often sacrificing one for the other.

Likewise, environmental impact studies have to balance conservation with the industry's interest in externalising as many costs as possible. Typically, a study has to show that the planned project meets environmental standards set by relevant authorities in terms of environmental indicators (e.g., water purity) and risk mitigation of catastrophic environmental incidents (e.g., oil spills). Because mitigation of environmental impact is expensive, companies have large incentives to aim towards the minimums required by the standards. Therefore, as in the case of activities in protected areas, environmental impact of industry is effectively managed towards the maximum acceptable environmental degradation.

Defining acceptability is not exclusively in the domain of environmental

sciences. Lower environmental standards attract investments, thus incentivising legislators to be as lax with the standards as the local population will allow. The willingness of the locals to tolerate environmental destruction partly depends on their awareness of environmental and health issues, but more so on their standards of living and needs for employment. Therefore, the balance between conservation and environmental degradation is set at the intersection of environmental sciences, economics, social sciences, health sciences, and politics. In countries where the population can afford to worry about environmental and health issues, red lines are drawn at ecological and health tipping points, that is, at environmental damage levels that, if increased, are all but guaranteed to cause unacceptable irreparable long-term damage to ecosystems or have clearly measurable effects on human health.

11.3. Tipping points

Tipping point is a point in parameter space at which a small perturbation induces a significant change in the state of the system (Fig. 68). Such points are a common feature of complex systems, and affect ecology at all scales—from molecular through organismal, individual and population, to ecosystem level.

Molecular-level tipping points in ecology typically occur when a negative feedback maintaining homeostasis is overcome by a forcing variable such as exposure to toxicants (Fig. 69). For low levels of exposure, the cell is able to maintain homeostasis by up-regulating molecular defence and repair mechanisms, even for large amounts of cellular damage. At the exposure level that exceeds the capacity of biomolecular control, even a small additional exposure will initiate a positive feedback loop of damage creation; additional cellular damage reduces the cell's ability to defend against exposure and increase in damage, thus accelerating (runaway) damage accumulation. Therefore, runaway damage accumulation initiated by a molecular tipping point (breakdown of molecular control) leads to death of the individual cell or organism, thus creating an individual-level tipping point. For complex systems, there could be intermediate states in which, for example, the organism is able to prevent runaway damage creation if damage levels are low, but not if the initial damage levels are high (Fig. 69). In such cases, any preexisting damage or additional damage-creating stress would result in death from exposure levels that would not harm an otherwise unstressed organism. This is particularly important to note when using experimental exposure data to interpret real-world situations in which organisms are rarely subject to a single




Figure 68: Example of a tipping point in ecosystems. Curves show an ecosystem’s equilibrium state as a function of environmental forcings such as nutrient abundance, temperature, salinity, precipitation, human exploitation etc. Arrows show the direction in which the ecosystem’s non-equilibrium state shifts over time given the conditions. **A**, Ecosystem equilibrium is a monotonic, slightly convex function of conditions. If the ecosystem state is initially below (above) the equilibrium curve, an upward (downward) shift is to be expected over time. **B**, Non-linear feedbacks may lead to pitchfork bifurcation where, for a given condition, multiple equilibria exist (here, three equilibria exist between points P and P'). This transforms the ecosystem’s convergence dramatically. If conditions in the tipping point P' change ever so slightly, the ecosystem undergoes a large forward shift. Even if conditions remain constant, a small perturbation close to the tipping point P' could induce the shift. This is because the unstable part of the equilibrium curve (dashed) indicates unstable equilibria around which the direction of convergence changes abruptly. Finally, akin to hysteresis in ferromagnetic materials, recovery following a forward shift at P' requires return of the forcing variable (conditions) all the way to the tipping point P . See Refs. [1040, 1041, 803, 802] for further details.

source of stress.

Food abundance is a major forcing in ecosystems that can initiate a number of ecological tipping points. For example, when food abundance reduces below corresponding tipping-point levels, ontogeny changes dramatically as individuals cannot grow, mature, or survive. Interestingly, however, even if an individual seems to be thriving, a population might have crossed a tipping point leading to species extinction (Fig. 70).

There is ample evidence of ecological tipping points, that is, tipping points at the level of whole ecosystems. Among the best-known examples is overfishing-induced ecosystem regime shift off the coast of Newfoundland in the early 1990s [1044, 1045]. The northern cod fishery that operated in the area was a product of systematic cod exploitation from the 16th cen-

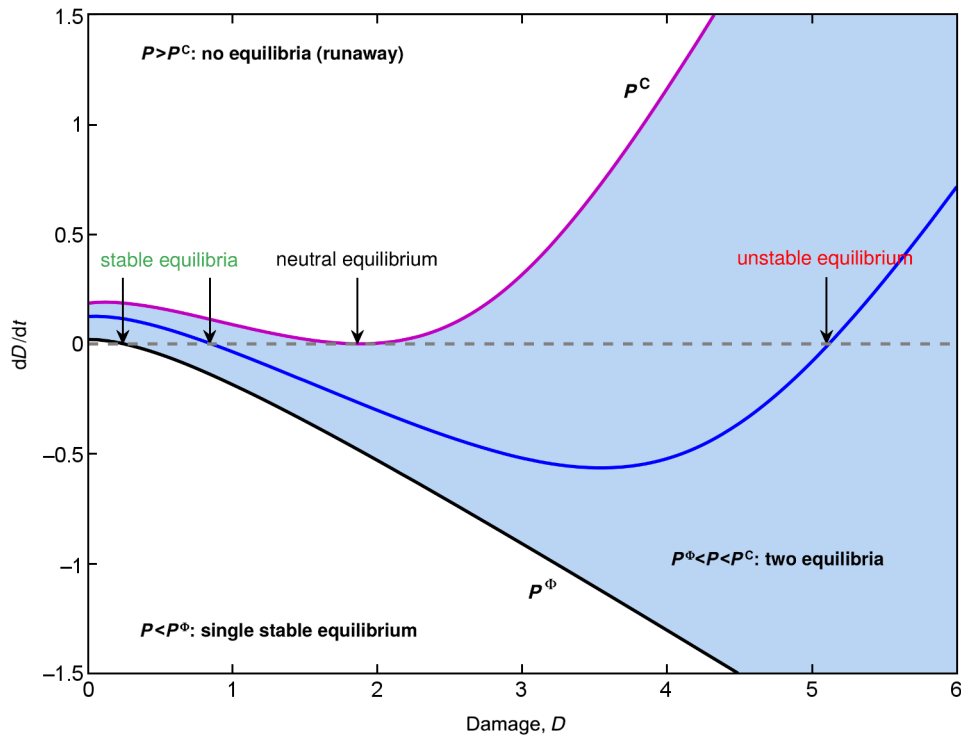


Figure 69: Example of a molecular tipping point turning to an individual-level tipping point. Solid curves delineate areas of distinct damage dynamics ($\frac{dD}{dt}$) as a function of damage (D); the dashed line represents an equilibrium ($\frac{dD}{dt} = 0$). As damage production due to exposure (P) increases from zero, the system has a single stable equilibrium, and damage levels are small. When P reaches a critical value of P^Φ (black solid curve), further increase in exposure results in two distinct equilibria, a stable and an unstable one (shaded area with a blue solid curve for illustration); damage levels are controlled unless initial damage levels are too high (D greater than the unstable equilibrium point for the given P). Further increase of P past P^C leads to a saddle-node bifurcation in that the two equilibria combine into one neutral equilibrium and disappear, thus initiating runaway dynamics for any D ; damage is not controlled, and the organism eventually dies from exposure. See Ref. [1042] for further details.

tury onwards. Since the 1880s, the fishery yielded over 200,000 tonnes of cod annually, but introduction of new technologies in the 1960s led to an explosion in yields which peaked at 810,000 tonnes in 1968. Technologies in question included more powerful trawlers equipped with radars, electronic navigation, and sonars. This enabled fishing longer, over larger areas, and at greater depths. Considerable amounts of non-commercial bycatch was

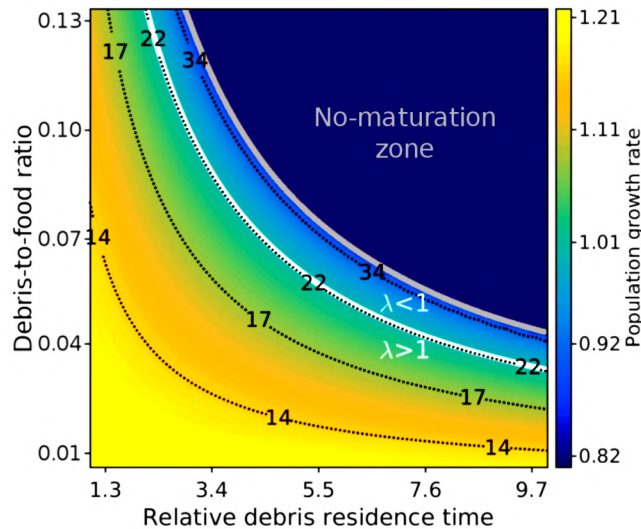


Figure 70: Individual- and population-level tipping points caused by plastics. Effects of plastic debris on loggerhead turtles were investigated using DEB. Black curves represent maturation age, while the colour on the graph represents population growth rate depending on the ratio of debris to food (y -axis), and residence time of debris in the gut (x -axis). If both are high enough, individuals can neither mature nor reproduce (the dark blue zone; individual-level tipping point). In this zone one would not observe reproducing individuals at all. If the ratio and the residence time are low enough, population is growing, and reproducing individuals can be observed. However, there is a zone (light blue between the white and grey curves) in which healthy reproducing individuals are observed, but the population is nevertheless going extinct (population-level tipping point). For loggerhead turtles this is a realistic prospect given that the ratio of debris to food required has already been observed. Process-based modelling expedites identification of such subtle but extremely important tipping point. Source: Reprinted figure from Ref. [1043].

extracted from the sea, including cod's prey, the capelin.

The marine ecosystem off the coast of Newfoundland finally gave in to enormous pressure in the early 1990s. In 1991, the fishery still landed about 129,000 tonnes of cod, which was roughly 69% of the plan for 1992. This plan, however, was entirely unrealistic, and in the same year, authorities announced a two-year moratorium on cod fishing in response to a dire state of fish stocks [1046]. The moratorium was eventually extended to a full decade, but even in 2002, there was no sign of cod recovery. Instead the ecosystem was dominated by invertebrates, crabs and shrimps.

First signs of recovery were reported in 2011 [1047], suggesting that prey-

fish stocks (such as the aforementioned capelin) had exploded in the wake of the cod-fishery collapse. This in turn had put enormous pressure on cod eggs and larvae. By 2005, however, prey-fish stocks themselves went into a downward spiral, providing a window of opportunity for the cod. An optimistic assessment of the situation was reiterated in 2015 [1048]. The reported recovery proposes a dynamically intriguing possibility; namely, a predator species (cod) historically held a prey species (capelin) in check, but upon the collapse of the predator, the prey multiplied enough to exert pressure on the predator's offspring, thus preventing a recovery. Such dynamics will be examined next through the prism of mathematical modelling.

Ref. [1049] presents a food-chain model consisting of an unstructured predator population feeding on a structured prey species, which in turn feeds on an unstructured resource (i.e., zooplankton) species. The term 'structured' designates that in the prey population there is an explicit distinction between juvenile and adult individuals. This is not the case in 'unstructured' populations. In the model, the ontogeny of prey individuals is specified in terms of three functions dependent on the resource abundance, R , and the individual's body size, l . *The ingestion function* is

$$I(R, l) = I_m l^2 \frac{R}{R_h + R}, \quad (119)$$

where I_m is the maximum surface-area-specific mass ingestion rate and R_h is the half-saturation constant, that is, $I(R_h, l) = \frac{1}{2} I_m l^2$. *The growth function* is

$$g(R, l) = \gamma \left[\frac{I(R, l)}{I_m l^2} l_m - l \right], \quad (120)$$

where γ is the maximum growth rate and l_m is the maximum body size achieved when the resource is abundant, that is, $R \rightarrow \infty$. *The reproduction function* is

$$b(R, l) = \begin{cases} 0, & \text{for } l < l_j \\ r_m \frac{I(R, l)}{I_m}, & \text{for } l \geq l_j \end{cases}, \quad (121)$$

where l_j is the length at sexual maturation and r_m is the maximum surface-area-specific reproduction rate. Additionally, the prey species is subject to a natural-mortality rate μ and a predator-induced mortality

$$d(P) = \begin{cases} 0, & \text{for } l < l_b \\ \frac{aP}{1+aT_h B}, & \text{for } l_b \leq l < l_v \\ 0, & \text{for } l \geq l_v \end{cases}, \quad (122)$$

where l_b and l_v specify the range of body sizes in which the prey species is vulnerable to predation, a is the maximum predator ingestion rate, and T_h is the mass-specific prey-processing time of the predator. The quantity P stands for the predator abundance, whereas the quantity B is the prey biomass accounting only for individuals with body sizes in the predation-vulnerability range, $l_b \leq l \leq l_v$.

The biomass B is not a state variable directly tracked by the model. Instead, the model keeps track of the prey density $c(t, l)$ of body size l , from which we have

$$B(t) = \int_{l_b}^{l_v} \beta l^3 c(t, l) dl, \quad (123)$$

where a typical weight-length relationship is used with the proportionality constant β . The dynamics of the state variable $c(t, l)$ is given in terms of a partial differential equation

$$\frac{\partial c(t, l)}{\partial t} + \frac{\partial}{\partial t} [g(R, l)c(t, l)] = -[\mu + d(P)] c(t, l). \quad (124a)$$

To be solvable, this equation needs a boundary condition, which is given in terms of the reproduction function

$$g(R, l)c(t, l) = \int_{l_j}^{l_m} b(R, l)c(t, l) dl. \quad (124b)$$

The dynamics of the state variable R is given in terms of an integro-differential equation

$$\frac{dR}{dt} = \rho(K - R) - \int_{l_b}^{l_m} I(R, l)c(t, l) dl, \quad (124c)$$

where ρ is the zooplankton inflow rate and K is the corresponding carrying capacity. The model is fully specified with a differential equation describing the predator dynamics

$$\frac{dP}{dt} = \epsilon d(P)B - \delta P, \quad (124d)$$

where ϵ is the growth efficiency of predator on prey and δ is the predator's natural mortality.




Figure 71: Zooplankton, prey, and predator dynamics in a marine ecosystem subjected to human exploitation of the predator species. Below the invasion threshold, zooplankton, adult prey, and predator are abundant, while juvenile prey is rare. Above the persistence threshold zooplankton and adult prey are rare, juvenile prey is abundant, while predator goes extinct. Between the two thresholds zooplankton and juvenile prey are inversely related, while adult prey will be abundant if predator is abundant too or adult prey will be rare if predator goes extinct. Red vertical lines (black curves) denote stable equilibria without (with) predator. Red arrows indicate the direction of the model's convergence, whereas the black arrow represents a perturbation affecting the abundance of adult prey. See Ref. [1049] for further details.

Source: Courtesy of Andre M. de Roos.

The described model generates extremely rich dynamics. Here, we are interested in the case of increasing human exploitation of the predator species, which increases the predator's mortality rate. Depending on this mortality rate, there are two key thresholds, the invasion threshold and the persistence threshold (Fig. 71). Below the former threshold, predator is abundant, as are adult prey and zooplankton, while juvenile prey is rare. Between the two thresholds, the predator abundance decreases with the mortality rate, allowing adult prey to reach its peak, which in turn means more offspring (i.e., juvenile prey) and a consequent zooplankton reduction. Above the latter threshold, however, predator suddenly goes extinct, adult prey and zooplankton become rare, while juvenile prey abounds. This is the ecosystem's tipping point.

Close to the tipping point various perturbations may turn the ecosystem's state upside down even if the predator mortality rate is below the persistence

threshold. For example, an adverse event affecting adult prey would be sufficient to change the model's convergence in such a way that adult prey becomes rare and predator ultimately collapses (Fig. 71). As a consequence, the ecosystem would be dominated by juvenile prey, just as was the case off the coast of Newfoundland where capelin and other prey-fish stocks became abundant after the cod fisher had gone bust.

11.4. *Future outlook*

Arguably, the primary socially relevant objective of modern environmental sciences should be to identify and improve our understanding of tipping points where irreparable environmental damage becomes inevitable. Indeed, incorporation of physics into environmental sciences leads to an unprecedented improvement in our understanding of such tipping points. Simultaneously, incorporation of economics through the concept of ecosystem services has raised awareness of trade-offs between short-term monetary gain from allowing externalisation of environmental costs, and long-term losses from reduced services provided by the ecosystem [1050]. Now more than ever, therefore, environmental sciences can help the decision-making process by (i) calculating positive and negative externalities, (ii) providing easily available large-scale longitudinal monitoring, and (iii) conducting risk assessment of ecological scenarios. Three major issues remain, though, implying new environmental research and management opportunities.

The mindset of acceptable change. Thinking that acceptable impacts do not have negative effects is referred to as the mindset of acceptable change. For example, even process-based models in ecotoxicology rely on the concept of no-effect concentration. Such a concentration, however, does not actually exist—even traces of a toxicant elicit responses at the biomolecular level, and can up-regulate a number of cellular processes [1051]. We do not observe effects of small concentrations on ontogenetic endpoints (e.g., reproduction or growth) only because of homeostatic regulatory mechanisms. These regulatory mechanisms are complex dynamical systems that utilise feedback loops to buffer against negative change, and maintain cellular or organismal homeostasis. Hence, detectable effect on an ontogenetic endpoint can be interpreted as a failure of the homeostatic buffering mechanism, and the no-effect concentration can be viewed as the concentration at which at least one homeostatic mechanism is about to fail. Is such a concentration safe, or can the up-regulation of the homeostatic mechanisms lead to harm in the

long term? Does the homeostatic mechanism maintain its effectiveness when environmental conditions change? Answering such questions requires models of toxic effects that incorporate regulatory dynamics at the organismal level, but we are aware of only one such model [1042]. The same considerations are valid when factoring acceptable environmental impact of human activities relevant in environmental impact studies and protected area management.

As discussed in the previous section, the mindset of acceptable change tends to result in management towards maximum environmental impact shy of tipping the environment into unwanted states. Perhaps it is therefore not surprising that, guided by funding, environmental sciences have been researching the overarching question of the acceptable change mindset: “How far can we stretch the environment?” We propose to change the question into something more constructive.

For example, in national park management, the question could be: “How can we improve the environment?” Surprisingly, the answer is not to ban tourism. National parks require expensive monitoring, management, scientific research, and active conservation efforts. Tourism provides the necessary funds either directly through entrance fees, or indirectly through local taxes collected from economic activities related to tourism. Hence, some tourism activity may be necessary for preservation, especially in poor regions where funding conservation is not high on the list of priorities. Tourism then contributes to the ability of a region to fund conservation but, even more importantly, employment opportunities in tourism provide motivation for conservation at all societal levels. Therefore, the new question does not affect whether tourists should be allowed to visit a protected area; it may, however, seek to limit the number of tourists, and affect their spatio-temporal distribution.

The difference in environmental impact between the two approaches can be huge. For example, while still in the mindset of acceptable change, two major national parks in Croatia (National Park Plitvice and National Park Krka) used to fight overcrowding by spreading the visitors throughout the park. This approach resulted in newly cleared paths, further habitat fragmentation, new conflicts between visitors and wildlife, and other negative environmental impacts. Due to a number of research projects, the mindset changed. Now, the two parks (i) streamline visitor experience to reduce perception of crowding, (ii) use monetary incentives to reduce peaks in demand, (iii) offer tourist attractions outside of the protected area, and (iv) limit simultaneous number of visitors in the area to safeguard visitor experi-

ence. These approaches enabled the parks to concentrate a greater number of tourists into a smaller area without sacrificing visitor experience. The higher concentrations of tourists made advanced environmental impact mitigation measures practical, thus drastically reducing the impact per visitor.

Re-visiting the management goals, therefore, reduced environmental degradation whilst increasing the number of visitors, their happiness, and income from entrance fees. Getting to the solutions, however, required interdisciplinary research and collaboration between social science, environmental science, and economics. We suggest goals of environmental impact studies, and environmental standards, should be revisited similarly.

The Big Picture. Ecosystems, responsible for the very air we breathe, depend on photosynthesis—processes on atomic levels that turn light into organic matter (but see Ref. [1052]). This leads to a local decrease in entropy essentially fuelled by nuclear reactions in the sun. All further interactions increase local entropy, hence our biosphere is limited by that initial step. Transferring the organic matter through the food chain is extremely inefficient; about 90% of energy is lost at each step [1053], thus the number of levels in the food chain is also limited. Therefore, as humans co-opt an increasing proportion of space and the photosynthetic production, ecosystems are put under increasing strain. Because ecosystems are complex dynamic systems, long-term consequences of changing the forcing of the system (and directly influencing all of its components) is extremely hard to predict.

Observations can provide only limited insights into small components of the ecosystem. Even large-scale monitoring can inform us about the past and the present, but not about the future. We can try to extrapolate from experiments, but they are of limited scope by definition, and therefore applicable to a limited set of environmental conditions and interactions. Furthermore, investigating organisms informs us about the individual, not the population over multiple generations in an uncontrolled environment.

To overcome the limitations of observation, we need to continue developing mechanistic models. Ideally, models would link gene expression to organismal ontogeny, to population status, to ecosystem dynamics. Then, we could truly explore optimal solutions to environmental and a host of societal and economic problems. Modelling must, however, be tempered by reality.

Virtual reality. Tools for linking sub-cellular processes to ecosystem-scale effects are being developed at an accelerating rate. Further development of

these tools will create an unprecedented *in-silico* test-bed for risk assessment and testing of potential ecological scenarios to minimise effects of anthropogenic pressures on the environment. However, increasing complexity of the underlying theory and physics in these tools makes it difficult for any one person to understand the background and assumptions of all the linkages. Simultaneously, the tools are increasingly easy to use, and their interfaces require less understanding of the underlying processes. The increasing complexity and ease-of-use could then combine to yield 'virtual reality' quasi-scientific consensus where model outputs drive policy even when divorced from reality. This effect can already be seen in fisheries where model-driven quota settings repeatedly fail to attain predicted results, and yet neither the procedures nor the modelling have been affected.

Dangers of succumbing to the virtual reality should be avoided by aggressively pursuing integration of data into modelling. Fortunately, relevant data collection is increasing exponentially. Biophysics generates knowledge on the (sub-)molecular level; remote sensing through satellites and drones, in combination with in-situ sensors creates snapshots of the environment going back decades; citizen science initiatives and obligatory reporting required by western governments contains a wealth of information on the environment. These sources are, however, not cross-linked and integrated, nor will they be without sufficiently accurate underlying models spanning relevant levels of biological organisation. Environmental science has only recently started the iterative cycle of mutually improving data collection and modelling, a process that brought physics to the forefront of natural sciences. The transition may require a change in attitude as well as competencies.

Ultimately, long-term solutions to environmental problems will be found at the intersection of environmental sciences, social sciences, economics, and politics. Environmental science can only identify problems and offer potential solutions. Choosing and implementing those solutions always has been a matter of policy and governance, which require social consensus; getting to the consensus is for social sciences to (re)solve. Economics clearly affects the social dialog, and ultimately makes implementation of solutions possible by finding ways to make environmentally sound choices in line with monetary incentives of the modern society. The high degree of ecosystem complexity implies that the best our civilisation can currently do is to aim at precautionary adaptive management fuelled by rigorous interaction between modelling and data. To achieve the needed breakthrough, a new generation of environmental scientists versed in mathematics, machine learning, and physics may

be required.

12. Global climate change

Climate change is among the greatest challenges that humanity has to face. Overwhelming scientific evidence indicates that a changing climate has tremendous influence on societies, both past and present, with serious consequences for the future [1054]. Recent advances in quantitative empirical research have illuminated the key connections in the coupled climate-human system [1055]. Numerous statistical analyses have addressed the causal effects between specific climatic conditions and social outcomes, such as agriculture, economics, conflict, migration, and health [1056]. Fig. 72 highlights a number of empirical studies that demonstrate how climatological events affect various social outcomes on both regional and global scales.

Many societal effects of climate have been evidenced. For example, temperatures around the 75th percentile of available records in Italy are associated with the lowest mortality there [1057]. An increase of one standard deviation in high-temperature days in India increases annual mortality among rural populations by 7.3 % [1058]. More intense tropical cyclones lead to more destruction in a global cross-section of countries [1059]. Non-linear temperature effects signal severe damages to the U.S. maize yields [1061]. Insufficient rainfall in Brazil causes steep drops in agricultural income [1062]. Labour productivity in the U.S. is strongly determined by the daily average temperature [1063, 1064], as is residential electricity consumption in California [1065], total income per capita in the U.S. [1067], and total factor productivity in China [1068]. The annual average temperature, furthermore, plays a key role in the growth of gross domestic product per capita [1066]. Interpersonal aggression, both petty and criminal, increases with temperature and sometimes decreases with rainfall; examples include the use of profanity on social media [1069] on the one end, and rape [1070] on the other end. Even civil conflicts escalate in the tropics in response to El Niño-type warming that takes place in the tropical central and eastern Pacific Ocean [1071]. Finally, changes in the multi-year average temperature have a greater effect on permanent outward migration of households in Indonesia than other natural disasters [1072]. All these quantitative empirical examples reveal that climate is indeed a major factor affecting social outcomes, often with first-order consequences. Even more important, however, is that understanding the relationship between climate and society offers insights into how mod-

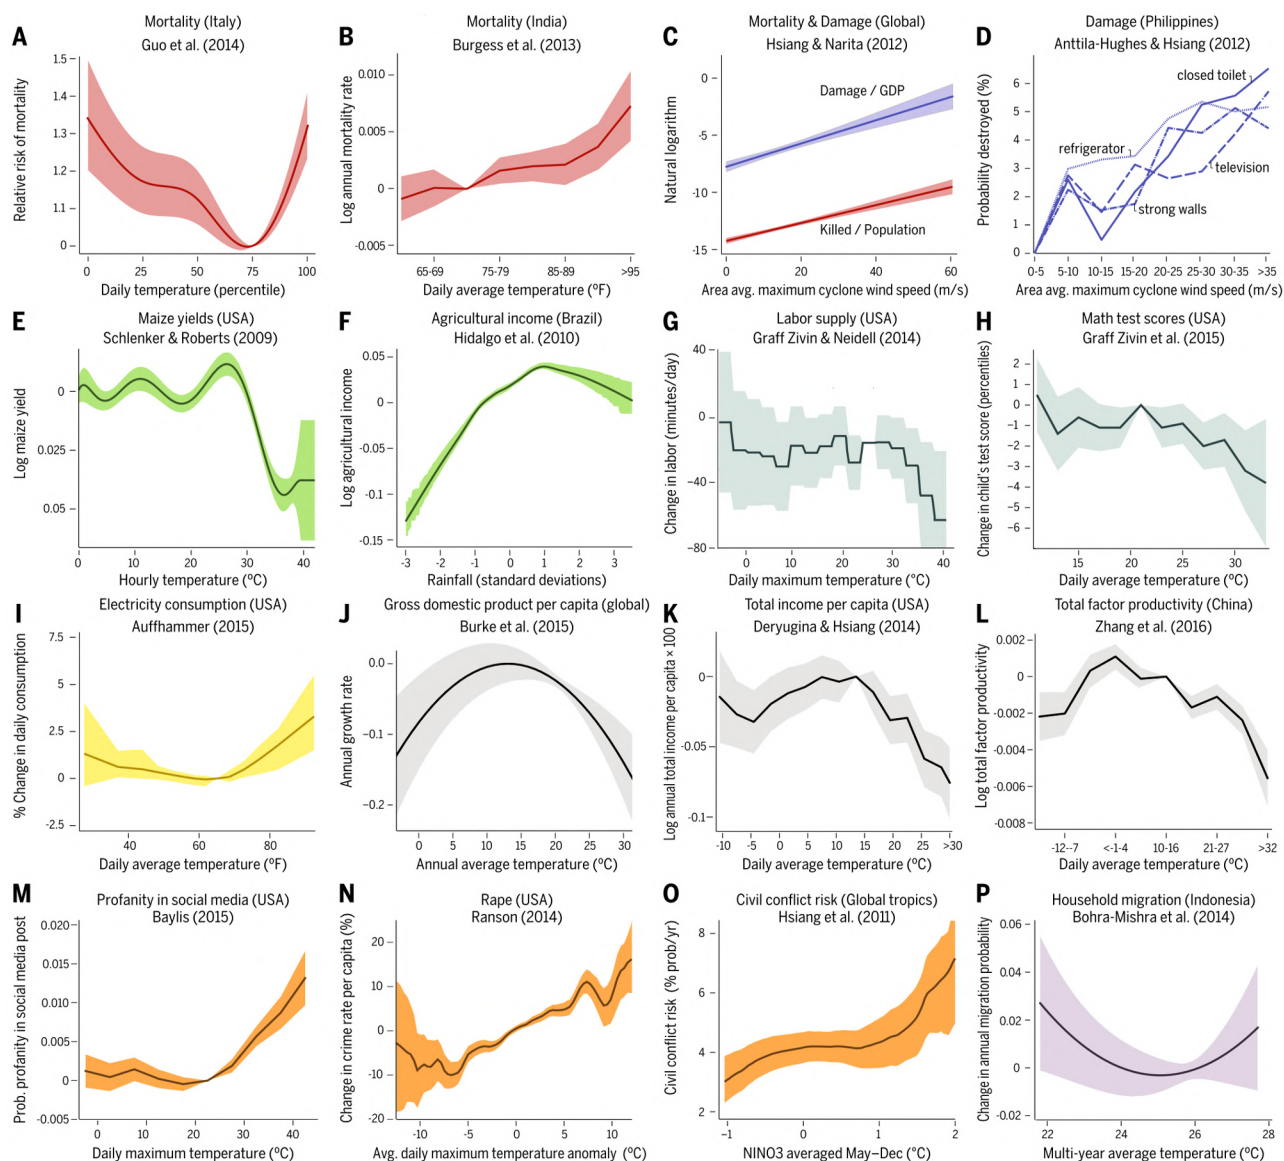


Figure 72: Social consequences of climate variables. The causal effect of climatological events on various social outcomes is described by a dose-response function. Colours indicate categories of outcome variables: **A, B**, red, mortality [1057, 1058]; **C, D**, blue, cyclone damage [1059, 1060]; **E, F**, green, agriculture [1061, 1062]; **G, H**, teal, labour productivity [1063, 1064]; **I**, yellow, electricity consumption [1065]; **J–L**, grey, aggregate economic indicators [1066, 1067, 1068]; **M–O**, orange, aggression, violence, and conflict [1069, 1070, 1071]; **P**, purple, migration [1072]. Shaded areas are confidence intervals. Source: Reprinted figure from Ref. [1056].

ern society can best respond to the current climatic events, and how future climate trajectories may impact humanity.

When considering the climate-human relationship, a crucial realisation is that, while climate drives social outcomes, human activities and emissions also impact the dynamics of climate change. According to the Intergovernmental Panel on Climate Change's (IPCC) fifth assessment report:

It is extremely likely that more than half of the observed increase in global average surface temperature from 1951 to 2010 was caused by the anthropogenic increase in greenhouse gases concentrations and other anthropogenic forcings together.

The underlying mechanism of human contributions to global warming is clear. The combustion of fossil fuels like coal and oil emits into the atmosphere greenhouse gases (GHGs), primarily carbon dioxide CO₂. GHGs block convective heat from escaping into space, which ultimately manifests as the rise in temperature.

Given the interdependence between climate and society, mitigating risks due to climate change requires an integrated perspective that goes beyond just the grasp of physical facts. It is also necessary to mobilise human action, which is, as a research problem, in the domain of multiple disciplines such as behavioural economics, social psychology, and evolutionary game theory. Scientists are attempting to rise to the challenge through the ongoing integration of climate science, social sciences, and humanities, giving rise to a new 'science of the Earth' called *Earth System Science* [1073]. The aim is to build a unified comprehension of the Earth and its human population.

Climate extreme events and global warming. Among the most visible consequences of climate change are increases in the intensity and frequency of extreme weather and climate events. These include heat waves, droughts, wildfires, floods, and hurricanes, to name a few. Such extreme events endanger not only human lives, but livelihood as well, as evidenced by, for example, fresh-water shortages and reduced food production. An extreme event is said to occur when the value of a climatic variable moves beyond the corresponding critical threshold. Of note, however, is that extreme events may be due to natural climate variability that is unrelated to anthropogenic forcings.

Extreme weather and climate events are grouped into three categories [1074]:
(i) extremes of atmospheric weather and climate variables, including tem-

perature, precipitation, and wind; (ii) weather and climate phenomena that influence the occurrence of extremes in weather or climatic variables, or represent extremes themselves, including monsoons, tropical cyclones, and extratropical cyclones; (iii) impacts on the natural physical environment, such as the aforementioned heat waves, droughts, floods and more. The distinction between these categories, while intuitive, is also somewhat blurred, and the categories are highly correlated.

The changes in the frequency, intensity, spatial extent, duration, and timing of weather and climate extremes may precipitate unprecedented risks and disasters for both natural physical environment and human society. Strengthening resilience against disruptive weather phenomena and climate change at national, regional, and local levels is therefore of vital importance. The ability to anticipate and predict extreme events would greatly aid the efforts to strengthen the resilience of human communities, but non-linear feedbacks, coupled interactions, and complex structure of the climate system pose formidable challenges. The situation is even more dire when attempts are made to account for the coupled climate-human system.

Tipping points in climate and social systems. The concept of a tipping point commonly refers to “a critical threshold at which a tiny perturbation can qualitatively alter the state or development of a system” [1075]. The term ‘tipping point’ was popularised by the writer Malcolm Gladwell who used it to describe intriguing sociological events during which little things can make a big difference [1076]. Many complex systems, including climate [1077, 1078, 1079] and human society [1080, 1081], have tipping points at which an abrupt shift to a contrasting dynamical regime may occur. In relation to the climate system, Ref. [1075] introduced the term ‘tipping element’ to describe large-scale components of the climate system that may approach or exceed a tipping point. Some examples are the ice-loss acceleration in Greenland, the diminishing sea-ice area in the Arctic, novel pests and fire patterns in boreal forest, the slowdown of the Atlantic thermohaline circulation, intense droughts in the Amazon rainforest, the large-scale die-offs of coral reefs, the decay of the Antarctic ice sheet, and others. Climate scientists have long suspected that by the present time, many tipping points of the climate system will have been exceeded, pushing the Earth ever closer to a *global tipping point*. Exceeding the global tipping point would constitute a downright existential threat to civilisation [1082]. It is, therefore, high time for international action, such as reducing the emissions of GHGs, to improve the planet’s resilience against

extreme climate.

Recently, Ref. [1083] proposed a framework of social tipping dynamics for stabilising Earth’s climate such that the planet is put back on track to halve global emissions by 2030 and tip the scales to net zero emissions by 2050. The study extends the idea of tipping elements from the components of the climate system to the subdomains of the planetary socioeconomic system, calling such subdomains *social tipping elements*, “where the required disruptive change may take place and lead to a sufficiently fast reduction in anthropogenic greenhouse gas emissions”. In this context, *social tipping interventions* have the potential to set social tipping elements on the path of change, for example, by (i) highlighting the moral implications of fossil fuels, (ii) strengthening climate education and engagement, and (iii) disclosing information on greenhouse gas emissions. By doing so, social tipping dynamics could be harnessed to foster climate-change mitigation.

Anticipating and predicting tipping points before they are breached would yield substantial socioeconomic benefits. Many techniques have been developed to this end based on the theory of early-warning signals for critical transitions [802]. The phenomenon of *critical slowing down* is considered as one of the most important clues that a dynamic system has lost resilience and is fast-approaching a tipping point [782]. Critical slowing down is recognised by an increase in the auto-correlation and the variance of the system’s state variables. For example, Ref. [1078] analysed eight abrupt shifts in ancient climate and found that significant increases in the lag-1 auto-correlation had preceded these shifts. Ref. [1084] examined a lake model in the vicinity of a bifurcation point and found an increasing variance of lake-water phosphorus about a decade prior to the shift to a new, nutrient-rich state. Other quantities and techniques have also been proposed as early-warning signals. Examples include the detrended fluctuation analysis [1085], power spectra [1086], flickering before transitions [1087], skewness and kurtosis [1088, 1089], and others. See Ref. [1090] for a more detailed review on the subject of tipping points and early-warning signals.

12.1. Modelling the climate system

What is a climate model? There are two main tools that supported the development of climate science: (i) observations of a changing Earth system and (ii) computer modelling and simulations. The term ‘observations’ in the context of climate science usually refers to instrumental data (i.e.,

meteorological stations or satellites), reanalyses data (e.g., ECMWF, NCEP-NCAR, and JRA), and proxy data (e.g., coral records, tree rings and ice-core records, etc.). ‘Computer modelling’ and climate models, by contrast, refer to attempts to simulate physical, chemical, and biological processes that take place in the atmosphere, cryosphere, land, ocean, and lithosphere and collectively produce climate. A climate model comprises a series of equations that describe said processes, and is typically implemented in numerical form that is suitable for processing on powerful computers. Crucially, scientists use climate models to project how climate may change over the course of the predictable future.

The history of climate modelling by means of numerical methods likely begins with Richardson’s work in the 1920s [1091], in which he proposed a novel idea to forecast weather using differential equations while viewing the atmosphere as a network of gridded cells. In 1938, Callendar published a seminal paper [1092] describing a one-dimensional radiative transfer model to show that rising CO₂ levels are warming the atmosphere. The first computerised, regional-weather forecast was tested in 1950 on the electronic numerical integrator and computer (ENIAC). The first three-dimensional general circulation model of the global atmosphere that could realistically depict seasonal patterns in the troposphere was released by Phillips in 1956. This was followed by the establishment of the National Center for Atmospheric Research (NCAR) in 1960, which soon thereafter became the leading climate modelling centre. The 1967 study by Manabe and Wetherald [1093] introduced an influential 1D radiative-convective model to generate the first credible prediction of the surface temperature in response to the CO₂ content of the atmosphere. NASA’s Nimbus III satellite was launched in 1969 with the specific task of taking measurements of Earth. The National Oceanic and Atmospheric Administration (NOAA) was created in 1970, and similar to NCAR soon thereafter became the world’s leading centre for climate-change research. The Met Office’s first general circulation model released in 1972 rounds up early developments in the field.

Rising awareness of climate change led to the establishment of the IPCC in 1988 with the aim to “provide the world with a clear scientific view on the current state of knowledge in climate change and its potential environmental and socio-economic impacts.” The first IPCC assessment report [1094] was published two years later with a summary stating that “under the IPCC Business-as-Usual emissions of greenhouse gases, the average rate of increase of global mean temperature during the next century is estimated to be about 0.3 °C

per decade.” With the more widespread development of coupled atmosphere-ocean global circulation models, a need arose for standardising their outputs, which resulted in the launch of Coupled Model Intercomparison Project in 1995. By the late 2000s climate models could be used in conjunction with paleoclimate data to explore climatic tipping elements [1075]. The biophysical understanding of Earth, including the climate system, was integrated with policy and governance in the planetary boundaries framework [1095]. Interestingly, between about 1998 and 2012, Earth seemed hardly to warm, which became known as the *global warming hiatus*, prompting some to question previous conclusions. The newer findings, however, reconciled models and data, leading the authors of Ref. [1096] to conclude that “we are now more confident than ever that human influence is dominant in long-term warming.”

A variety of climate models to date, from simple energy-balance models to elaborate general circulation models, differ in their complexity and relative advantages. Especially the general circulation models are highly reliant on the most advanced supercomputers in order to assimilate all the required data. Despite the rapid advancement of the field since its inception, there is no single, comprehensive model that could encapsulate all the non-linear interactions between climate-determining subsystems. Consequently, the longer-term predictive skills regarding climate variability and climate change remain limited and dependent on the precise initial conditions.

Hierarchy of climate models. Earth’s climate is a complex system influenced by many factors, for example, solar radiation, clouds, winds, ocean currents, and many others. The system is furthermore subdivided into subsystems—the atmosphere, the ocean, the cryosphere, the biosphere, the pedosphere, and the lithosphere—that interact at the interfaces such as air-ocean, air-ice, ice-ocean, as well as land-air and land-ocean. Over the years, climate modelling has benefited from a variety of approaches to constructing climate models that integrate, to a larger or lesser degree, said components and interactions. In particular, much has been learned from models focused on specific aspects of the climate system (e.g., El Niño events and monsoons), while abandoning the pretence that full complexity can be accounted for. This line of thinking and climate model development is now known as a *hierarchical modelling approach* [1097].

We distinguish four model categories based on their complexity:

1. *Energy-balance models* estimate the changes in the climate system by analysing Earth’s energy budget, that is, by balancing the incoming

solar radiation and the outgoing terrestrial radiation.

2. *Radiative-convective models* simulate the vertical profile of atmospheric temperature and the associated transfer of energy under the assumption of radiative-convective equilibrium.
3. *Statistical-dynamical models* combine the features of energy-balance and radiative-convective models in order to study horizontal energy flows and processes that disrupt such flows.
4. *General circulation models* attempt to capture the fundamental physics and chemistry of the climate system, including the exchange of energy and materials between the components of this system.

An energy-balance model can take one of two simple forms, the zero-dimensional model such that Earth is a single compartment with a global mean effective temperature or the one-dimensional model such that temperature is latitudinally resolved. In the one-dimensional model, each latitudinal zone is described by the following equation

$$(\text{Shortwave in}) = (\text{Transport out}) + (\text{Longwave out}), \quad (125)$$

or more formally

$$S(\phi)\{1 - \alpha(\phi)\} = c\{T(\phi) - \bar{T}\} + \{A + BT(\phi)\}, \quad (126)$$

where $S(\phi)$ is the mean annual radiation incident at latitude ϕ , $\alpha(\phi)$ is the albedo at latitude ϕ (0.62 for $T < -10^\circ\text{C}$ and 0.3 otherwise), c is the horizontal heat-transport coefficient ($3.81 \text{ W m}^{-2} \text{ }^\circ\text{C}^{-1}$), $T(\phi)$ is the surface temperature at latitude ϕ , \bar{T} stands for the mean global surface temperature, and A and B are constants governing the longwave radiation loss ($A = 204.0 \text{ W m}^{-2}$ and $B = 2.17 \text{ W m}^{-2} \text{ }^\circ\text{C}^{-1}$). Of note is that some implementations of energy-balance models also simulate energy transfers between the atmosphere and the ocean.

Radiative-convective models add complexity relative to the energy balance models. Thus, one-dimensional radiative-convective models account for the vertical dimension, while two-dimensional models additionally account for one horizontal dimension. Using such models it possible to predict how

GHGs modify effective emissivity and surface temperature. A radiative-convective model has the following mathematical form

$$S = \alpha_c S + \alpha_g (1 - a_c)^2 (1 - \alpha_c) S + \varepsilon \sigma T_c^4 + (1 - \varepsilon) \sigma T_g^4, \quad (127a)$$

$$a_c (1 - \alpha_c) S + a_c \alpha_g (1 - a_c) (1 - \alpha_c) S + \varepsilon \sigma T_g^4 = 2\varepsilon \sigma T_c^4, \quad (127b)$$

$$(1 - \alpha_g) (1 - a_c) (1 - \alpha_c) S + \varepsilon \sigma T_c^4 = \sigma T_g^4. \quad (127c)$$

These equations respectively represent the energy balances at (i) the top of the atmosphere, (ii) the cloud level, and (iii) the surface. The equations are directly solvable upon setting the values for the cloud shortwave absorption, a_c , the cloud albedo, α_c , infrared emissivity, ε , and the surface albedo, α_g . The parameter σ is the Stefan-Boltzmann constant. Radiative-convective models, as seen here, incorporate information about radiation fluxes throughout the atmosphere, including the fluxes of solar radiation, cloud cover, and land.

While statistical-dynamical models make more of a practical leap, general circulation models mark the next true conceptual leap in climate modelling. General circulation models are the most complex and ‘complete’ model type used in climate-change science. These three-dimensional models are constructed by discretising the differential equations that express the conservation of momentum, mass, and energy. The model closure is achieved by adding an equation of state for the atmosphere. Expressed in a mathematical form, we have:

1. *Conservation of momentum*

$$\frac{D\mathbf{v}}{Dt} = -2\boldsymbol{\Omega} \times \mathbf{v} - \rho^{-1} \nabla p + \mathbf{g} + \mathbf{F}. \quad (128)$$

2. *Conservation of mass*

$$\frac{D\rho}{Dt} = -\rho \nabla \cdot \mathbf{v} + C - E. \quad (129)$$

3. *Conservation of energy*

$$\frac{DI}{Dt} = -p \frac{D}{Dt} \rho^{-1} + Q. \quad (130)$$

4. *Ideal gas law*

$$p = \rho RT. \quad (131)$$

The physical meanings of the symbols are as follows:

- \mathbf{v} =velocity relative to Earth,
- t =time,
- $\frac{D}{Dt} = \frac{\partial}{\partial t} + \mathbf{v} \cdot \nabla$ =total time derivative,
- Ω =Earth’s angular velocity vector,
- ρ =atmospheric density,
- p =atmospheric pressure,
- \mathbf{g} =apparent gravitational acceleration,
- \mathbf{F} =force (other than gravity) per unit mass,
- C =creation rate of atmospheric constituents,
- E =destruction rate of atmospheric constituents,
- $I = c_p T$ =internal energy per unit mass,
- T =temperature,
- Q =heating rate per unit mass,
- R =gas constant, and
- c_p =specific heat of air at constant pressure.

Aside from the aforementioned model types, there are other classes of climate models that attempt to capture specific aspects of the climate system, but in a simplified way. These are known as *intermediate complexity models*. A representative example is the Cane-Zebiak model [1098], developed to simulate El Niño events, as well as conduct experimental climate predictions.

Notably, the physical models of the climate system lack any form of human dynamics, treating instead Earth’s global population as an outside force. Attempts to fill this gap produced integrated assessment models in which human dynamics plays a prominent part [1099, 1100]. These models aim to link socioeconomics with the biosphere and the atmosphere into one modelling framework for the purpose of simulating costs of specific climate-stabilisation policies.

Integrated assessment models “represent many of the most important interactions among technologies, relevant human systems (e.g., energy, agriculture, the economic system), and associated greenhouse gas emissions in a single integrated framework” [1054]. This means not only an integrated representation of the physical laws driving natural systems, but also the changing preferences that drive human society. Typically, there are two main types of integrated assessment models—simple and complex. Simple models are run in a spreadsheet by utilising simplified equations, while detailed relationships between the economy, energy, and Earth systems are left out [1101]. These models are commonly used to evaluate the ‘social cost of carbon’. By con-

trast, complex integrated assessment models account for energy technologies and uses, changes in land use, and societal trends. Separate modules represent the global economy and the climate system. The basic structure of an integrated assessment model can be broken down as follows:

- *Model inputs*—assumptions about how the world works and changes, such as the GDP, populations, policies, and so on;
- *Model itself*—modules that represent economic, energy, land, and climate systems;
- *Model outputs*—quantitative predictions about the economy, land-use changes, greenhouse gas emissions and energy-use pathways, and future human development.

Despite their comprehensiveness, complex models are imperfect, often failing to capture nuanced social mores, habits, and behaviours, both present and future. Such imperfections notwithstanding, complex integrated assessment models are a valuable tool in exploring the contributions of social factors to climate change. The key questions that can be addressed in this way include, for instance, how to avoid global warming of more than 1.5 °C at the lowest cost or what the implications are of current national pledges to reduce the emissions of GHGs. Taken together, global circulation, integrated assessment, and other climate models offer powerful means to explore Earth system dynamics at a range of spatial and temporal scales, all the while incorporating both physical and social mechanisms and processes.

12.2. *Climate networks*

In recent years, network science has emerged as a novel framework to study climate phenomena such as El Niño–Southern Oscillation (ENSO), extreme-rainfall patterns, and air-pollution variability [1102]. It is, of course, worthwhile to review this topic in its own right, but even more so given that network science interfaces physics with so many other disciplines.

Basic concepts. Networks have proven to be a versatile tool to explore the structural and dynamical properties of complex systems beyond physics, for example, in biological, ecological, and social sciences [485]. A particular strength of the network representation of a complex system is the ability to map out the system’s topological features. Climate, as mentioned previously, is a quintessential example of a complex system comprising many non-linearly coupled subsystems with multiple forcings and feedbacks. The desire to model climate’s complexity has, therefore, led to the birth of the idea

of *climate networks* in which geographical locations on a longitude-latitude grid become network nodes, while the degree of similarity, or ‘connectedness’, between the climate records obtained at two different locations determines whether there is a network link between these locations [1103, 1104]. The climate-network framework has been applied successfully to analyse, model, and predict various climate phenomena, such as ENSO [1105, 1106, 1107, 1108, 1109, 1110, 1111, 1112], extreme rainfall [1113, 1114], Indian summer monsoon [1115], Atlantic meridional overturning circulation [1116, 1117], Atlantic multidecadal oscillation [1118], teleconnection paths [1119], the impacts of CO₂ [1120], and others.

Climate networks are constructed and utilised in three steps (Fig. 73):

1. *Step 1.* Define a spatial grid of nodes containing the climatological variable of interest (e.g., temperature, geopotential height, precipitation, etc.).
2. *Step 2.* Build links between nodes pairs based on statistical correlations between the time series recorded at the two node locations.
3. *Step 3.* Interpret the dynamical processes of the climate system (e.g., winds, ocean currents, atmospheric circulation, Rossby waves, etc.) via the structural properties of the climate network.

A detailed overview of methodology for constructing and analysing climate networks can be found, for example, in Ref. [1121].

Building links between nodes is of central importance in constructing climate networks. Among the most direct ways to decide whether, or how strongly, two nodes are connected is the *Pearson correlation*. Let us suppose that a climate observable T (e.g., the sea-surface temperature anomaly) is measured at a number of fixed stations. At station i , which is to be identified with the i th node in the climate network, measuring the climate observable yields a time series $T_i(t)$. If the time series is subdivided into, for instance, calendar years, months, or days, this is further indexed with the index y producing $T_i^y(t)$. Then, the time-delayed Pearson cross-correlation function between nodes i and j is [1112],

$$C_{i,j}^y(-\tau) = \frac{\langle T_i^y(t)T_j^y(t-\tau) \rangle - \langle T_i^y(t) \rangle \langle T_j^y(t-\tau) \rangle}{\sqrt{\langle (T_i^y(t) - \langle T_i^y(t) \rangle)^2 \rangle} \cdot \sqrt{\langle (T_j^y(t-\tau) - \langle T_j^y(t-\tau) \rangle)^2 \rangle}}, \quad (132)$$




Figure 73: Constructing and utilising a climate network. In step 1, a spatial grid of nodes where climatological time series have been observed is defined. In step 2, the cross-correlation between time series at two node locations is calculated. When such a cross-correlation is strong, the node locations are deemed to be linked. In step 3, the topology of the climate network is analysed using the methods of network science to reveal the properties of the climate system.

and

$$C_{i,j}^y(\tau) = \frac{\langle T_i^y(t - \tau)T_j^y(t) \rangle - \langle T_i^y(t - \tau) \rangle \langle T_j^y(t) \rangle}{\sqrt{\langle (T_i^y(t - \tau) - \langle T_i^y(t - \tau) \rangle)^2 \rangle} \cdot \sqrt{\langle (T_j^y(t) - \langle T_j^y(t) \rangle)^2 \rangle}}, \quad (133)$$

where $0 \leq \tau \leq \tau_{\max}$ is the time lag and $\langle \cdot \rangle$ denotes averaging over the variable t . Next, it is possible to define the positive and negative link strengths of between nodes i and j as [1122]

$$W_{i,j}^{+,y} = \frac{\max(C_{i,j}^y) - \text{mean}(C_{i,j}^y)}{\text{std}(C_{i,j}^y)}, \quad (134)$$

and

$$W_{i,j}^{-,y} = \frac{\min(C_{i,j}^y) - \text{mean}(C_{i,j}^y)}{\text{std}(C_{i,j}^y)}, \quad (135)$$

where max, min, mean, and std respectively denote the maximum, minimum, mean, and the standard deviations of the cross-correlation function over the

time lag τ . The direction of the link is taken to be from node i to node j if $W_{i,j}^{+,y} > |W_{i,j}^{-,y}|$ and from node j to node i otherwise [1111].

An alternative method to construct climate networks is the *event synchronisation* method, which was originally developed to measure synchronisation and infer the direction of time delay between signals [1123]. The method is based on the relative timings of events in a time series, where an event is defined as, for instance, reaching a threshold or a local maximum. Supposing two time series, $X(t)$ and $Y(t)$, an event l seen in X at time t_l^x is considered to be synchronised with an event m seen in Y at time t_m^y , if $0 < |t_l^x - t_m^y| < \tau_{lm}^{xy}$, where

$$\tau_{lm}^{xy} = \frac{1}{2} \min \{t_{l+1}^x - t_l^x, t_l^x - t_{l-1}^x, t_{m+1}^y - t_m^y, t_m^y - t_{m-1}^y\}. \quad (136)$$

The quantity τ_{lm}^{xy} represents a minimum time lag between two consecutive events of the same type occurring in one of the two time series. Synchronisation thus requires that when the event of interest is seen in one time series, say, X , the same event should occur in the other time series Y too before being seen again in the time series X . If, furthermore, e_x and e_y denote the number of events in X and Y , respectively, then $l = 1, 2, \dots, e_x$ and $m = 1, 2, \dots, e_y$. A counter of instances when the event happens in X shortly after happening in Y is

$$c(x|y) = \sum_{l=1}^{e_x} \sum_{m=1}^{e_y} J_{xy}^{lm} \quad (137)$$

with

$$J_{xy}^{lm} = \begin{cases} 1, & \text{if } 0 < t_l^x - t_m^y \leq \tau_{lm}^{xy}, \\ 1/2, & \text{if } t_l^x = t_m^y, \\ 0, & \text{otherwise.} \end{cases} \quad (138)$$

Analogous reasoning is used to define $c(y|x)$. Finally, the symmetrical and anti-symmetrical combinations of these counters are

$$Q_{xy} = \frac{c(y|x) + c(x|y)}{\sqrt{e_x e_y}}, \quad q_{xy} = \frac{c(y|x) - c(x|y)}{\sqrt{e_x e_y}}. \quad (139)$$

Here, Q_{xy} measures the strength of event synchronisation, while q_{xy} estimates the direction of time delay. In the construction of climate networks, the former (latter) quantity determines the link strength (direction).

Other methods to quantify time-series, and thus node, similarity exist and can be used in the construction of climate networks. Examples are the *mutual information* method [1118] and the ϵ -*recurrence* method [1124]. Whichever method of quantifying similarity is adopted, it is common to discard weak links by applying a thresholding criterion. The network can, in fact, be made unweighted and undirected by defining the adjacency matrix as

$$A_{i,j} = \mathcal{H}(W_{i,j} - W_c). \quad (140)$$

where \mathcal{H} is the Heaviside function, $W_{i,j}$ are link weights calculated using, say, the Pearson cross-correlation, and W_c is a threshold. Once the climate network is constructed, it can be subjected to structural analyses using the methods of network science. Several examples of such analyses are outlined next, showing how the climate-network framework reveals new knowledge about climatic events of great societal relevance, for example, ENSO, extreme rainfall, and air pollution.

El Niño-Southern Oscillation (ENSO) forecasting. ENSO is among the most prominent phenomena of climate variability on the interannual time scale [1125, 1126]. The term refers to fluctuations between anomalous warm El Niño and cold La Niña conditions in the eastern Pacific Ocean. The occurrence of an El Niño event can trigger numerous disruptions around the globe, causing climate-related disasters, such as droughts, floods, fishery declines, famines, plagues, and even political and social unrest. To adequately prepare for these potential disruptions, it is pivotal to develop reliable prediction skills for when and where climate may turn extreme. After the first forecasting model from the 1980s, that is, the aforementioned Cane-Zebiak model, a number of dynamical and statistical models have been proposed to predict the El Niño events. International Research Institute for Climate and Society, for example, offers some 20 climate models for ENSO forecasts. Model richness notwithstanding, early and reliable ENSO forecasting remains a substantial challenge. Good prediction skill is generally limited to about 6 months ahead, due to the presence of the so-called ‘*spring predictability barrier*’, which greatly amplifies errors arising from the coupling and feedbacks in the equatorial atmosphere-ocean system [1127].

To improve the El Niño forecasting skill, especially beyond the spring predictability barrier, Ref. [1128] resorted to an approach based on climate networks that yields reliable predictions about one year in advance. Nodes

for the construction of the climate network were mainly located in the tropical Pacific (Fig. 74, upper panel), with only a minority of nodes, denoted in red, inside the El Niño basin. Link weights were calculated using the Pearson cross-correlation method (see Eqs. 132 and 134). To obtain the mean strength of dynamical teleconnections in the climate network, link strength was averaged across all links

$$W^y = \frac{1}{n_1 n_2} \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} W_{i,j}^{+,y}. \quad (141)$$

Here, n_1 and n_2 respectively stand for the number of red and blue nodes in the upper panel of Fig. 74. The quantity W^y was then compared with a decision threshold $\Theta = 2.82$, which had to be crossed from below in order to consider signalling an alarm. Additionally, the NINO3.4 index, that is, NOAA’s primary indicator for monitoring El Niño and La Niña events, needed to be below 0.5°C . If both conditions were satisfied, then the alarm would be signalled for the following calendar year (Fig. 74, lower left panel). The prediction accuracy of the climate network, as measured by a type of receiver-operating-characteristic analysis, turned out to be much higher than that of the state-of-the-art climate models, for example, Kirtman [1129] and Chen-Cane [1130] models. The approach, in fact, successfully predicted in 2013 the onset of the 2014–2016 strong El Niño event (Fig. 74C, lower right panel).

Afterwards, Ref. [1110] proposed another framework for predicting the onset of El Niño events that combined a time-evolving, weighted climate network with the elements of percolation theory. In this approach, nodes near-homogeneously covered the entire globe rather than just the tropical Pacific. The climate network was shown to undergo abrupt percolation transitions usually about one year before an El Niño event, thus providing a reliable early-warning indicator. These research efforts were followed by yet another approach based on climate networks that successfully predicted one year in advance the onset of the 2018–2019 El Niño event [1112]. In this last approach the climate network was located entirely in the El Niño basin.

ENSO greatly affects atmospheric circulation patterns and exhibits strong regional and remote influences on weather. To investigate the global impacts of ENSO, Ref. [1111] resorted to constructing a series of directed and weighted climate networks based on the near-surface air temperature. Regions that are characterised by larger positive or negative network links cor-




Figure 74: Climate network scheme for forecasting El Niño events. The upper panel shows the geographical locations of grid points (i.e., nodes). The network consists of 14 nodes in the El Niño basin (solid red symbols) and 193 nodes outside this domain (open symbols). The red rectangle denotes the NINO3.4 region (5°S–5°N, 170°W–120°W). In the lower left panel, the red curve is the average link weight W of the climate network as it changes through time, the horizontal line indicates the decision threshold $\Theta = 2.82$, and the blue areas show the El Niño events. When W crosses the threshold from below, an alarm is sounded indicating that there is an impending El Niño event in the following calendar year. Correct predictions are marked by solid green arrows and false alarms by dashed black arrows. The lower right panel is a magnification for August (A), September (S), October (O), and November (N) of 2013. Source: Reprinted figure from Ref. [1128].

related more with the NINO3.4 index, thus becoming warmer (cooler) during El Niño (La Niña) periods. Although regions affected by ENSO vary from one event to another, and are difficult to predict, the climate network analysis offered a new perspective on the problem with much potential for further

successes.

ENSO, as referred to heretofore, is sometimes called Eastern Pacific ENSO to distinguish it from a temperature anomaly that arises in the central Pacific [1131], which is called Central Pacific ENSO [1132] or ENSO Modoki [1133] (Japanese ‘modoki’ roughly translates as ‘pseudo’). ENSO Modoki has distinct teleconnections and affects many parts of the world, yet distinguishing and predicting the type of ENSO in practice remains a challenge. Climate networks may help, as evidenced by a novel method to predict the type of El Niño events, as well as estimate their impacts in advance [1134].

Extreme-precipitation patterns. Precipitation, at its extremes, poses a threat to society, resulting in the loss of life and property in floods and landslides. Flooding due to extreme precipitation in India, for example, affected over 800 million people in the period between 1950–2015, leaving 17 million without homes and causing 69,000 deaths [1135]. In early 2017 in coastal Peru, a series of extreme precipitation events caused severe floods, killing 114 people, displacing 184,000 people, and creating damages in excess of \$3 billion USD [1136].

Even more disconcerting than historical records is the fact that the intensity of extreme weather events is expected to strengthen under global warming. The basic mechanism of how a temperature rise fuels extreme rainfall is clear, (i) warmer ocean waters carry energy more easily to the atmosphere when storms form, and (ii) for every degree of surface-temperature warming, the atmosphere holds about 7% more water vapour [1137]. Accordingly, climate models also predict intensification in the annual maximum precipitation, although they possibly underestimate the true future state of affairs [1138], thus exposing gaps in our understanding of the factors involved. Such gaps, for example, include limited knowledge of global and regional teleconnection patterns associated with extreme rainfall. Recent progress in this context has relied heavily on climate networks, not only by mapping extreme-rainfall teleconnections, but also suggesting the underlying mechanisms behind the observed phenomena.

Ref. [1113] offered a new conceptual route to study the spatial characteristics of the synchronicity of extreme rainfall in South America during the monsoon seasons. First, the study defined extreme-rainfall events as those above the 99th percentile over the spatial domain covering 40°S–15°N and 30°W–85°W at a resolution of 0.25°, and the temporal domain extending from 1998 to 2012 at a resolution of 3 hours. A climate network was then

constructed using the aforementioned event-synchronisation method. The synchronisation strength into S_i^{in} (out of S_i^{out}) a climate-network node was defined as the sum of weights of all links pointing to (from) this node. The network divergence ΔS was then introduced to spatially resolve the temporal order of extreme-rainfall events,

$$\Delta S_i = S_i^{\text{in}} - S_i^{\text{out}} = \sum_{j=1}^N A_{ij} - \sum_{j=1}^N A_{ji}, \quad (142)$$

where A_{ij} is the adjacency matrix of the climate network. The positive (negative) values of ΔS_i indicated sink (source) nodes, that is, locations where extreme events occur shortly, within two days, after (before) occurring at many other locations. Typical propagation pathways of extreme events could thus be identified along which extreme events have high predictability. The method was applied to the real-time satellite-derived rainfall data to successfully predict more than 60 % of extreme-rainfall events in the Central Andes of South America, with the success rate going above 90 % during El Niño conditions.

A similar methodology based on climate networks, and specifically on the event-synchronisation method, has been applied to investigate the spatial configuration of synchronisation between extreme-rainfall events around the globe. Ref. [1114], using a network with 576,000 nodes, found that the distribution of distances between significant links (p-value $p < 0.005$) decays according to a power law with a coefficient ≈ 1 up to distances of about 2,500 km, while the probability of significant longer-distance links is much larger than expected from the power law. The relative underabundance of shorter-distance links is due to regional weather systems, yet the relative overabundance of longer-distance links, which form a global rainfall teleconnection pattern, is probably dominated by the *Rossby waves*. The described picture is robust to the choice of the extreme-event percentile (Fig. 75, left-column panels). Furthermore, climate networks revealed that the extreme-rainfall events in the monsoon systems of south-central Asia, east Asia, and Africa are strongly synchronised (Fig. 75, right-column panels). The use of climate networks thus made inroads towards the global predictability of natural hazards associated with extreme rainfall.

12.3. Impact of Rossby waves on air pollution

Air pollution is a major health concern worldwide. According to the World Health Organization (WHO) [1139]:




Figure 75: Distance distribution and teleconnection pattern in south-central Asia for different extreme-event percentiles. Left-column panels show the probability density function of the significant link distances (red and blue circles), the power-law fit over the range 100–2,500 km (dashed line), and the kernel-density estimate (KDE) of the distribution of all possible great-circle distances (solid black line) for extreme-event percentiles $\alpha = 0.94$, 0.95 , and 0.96 , respectively. Right-column panels show link bundles attached to south-central Asia for extreme-event percentiles $\alpha = 0.94$, 0.95 , and 0.96 , respectively. Links shorter (longer) than 2,500 km are denoted in red (blue). Source: Reprinted figure from Ref. [1114].

An estimated 9 out of 10 people worldwide are exposed to air pollutants that exceed World Health Organization (WHO) air quality guidelines. [P]olluted air kills some 7 million people each year, causes long-term health problems, such as asthma, and reduces children's cognitive development. According to the World Bank, air pollution costs societies more than \$5 trillion every year.

Climate change and air pollution are closely related. For example, the main sources of CO₂ emissions are also a major source of air pollutants. Conversely, many air pollutants, such as particulate matter, ozone, nitrogen dioxide, etc. contribute to climate change by affecting the amount of incoming radiation that is reflected or absorbed by the atmosphere. The exact feedbacks between weather and climate dynamics at different pressure levels, on the one hand, and the fluctuations in air pollution, on the other hand, is a subject of intense study.

To study air-pollution spreading and diffusion patterns, Ref. [1140] employed a multilayer and multivariable network analysis designed to delineate the influence of the upper air dynamics (at 500 hPa geopotential height) on the temporal variability of the surface air pollution (PM_{2.5}) in China and U.S. Two multilayer networks were considered, one with dominant negative-correlated interlinks and the other with positive-correlated interlinks, corresponding to negative, Eq. (135), and positive, Eq. (134), weights, respectively. Only the links for which $|W| > W_c = 4.5$ were selected based on shuffled data-significance tests. Applying this methodology showed that the upper air critical regimes substantially influence the surface air pollution. Specifically, Rossby waves influence the air-pollution fluctuations through the development of cyclone and anticyclone systems that control local winds and air stability (Fig. 76). High-pressure anticyclones form on the ridges, while low-pressure cyclones form on the troughs of Rossby waves. The former, identified by negative out-degree clusters in climate networks, cause the air to downwell. The latter, identified by positive out-degree clusters in climate networks, cause the air to upwell. The described downwelling and upwelling pattern induces strong winds that keep air pollution low. As Rossby waves travel, upwelling replaces downwelling and vice versa, which weakens the winds and leads to subsequent accumulation of pollution in the air. Recognising the outlined mechanism behind the air-pollution fluctuations helps to improve the prediction of extreme pollution events, and once again highlights the potential of climate networks to unveil intricate interac-




Figure 76: Schematic representation of Rossby waves influencing air pollution. Panel (a) shows low-pollution conditions. Panel (b) shows high-pollution conditions. Blue and red colours respectively represent the negative and positive out-degree clusters of the climate network.

Source: Reprinted figure from Ref. [1140].

tions and feedbacks in the climate system.

12.4. Critical phenomena in the climate system

Critical points, exponents, and universality. The concept of critical phenomena is most commonly associated with physical systems that undergo phase transitions at a *critical point*. Examples include the vapour-to-liquid-to-solid transitions of substances at their critical points characterised by a specific value of temperature and pressure, or the ferromagnetism-to-paramagnetism transition of some solids at their Curie point under zero magnetic field. Critical phenomena are present in both nature (lakes, oceans, terrestrial ecosys-

tems, etc.) and society [1141], including the climate system itself. Ref. [781] offers a general overview of the subject.

Curiously, critical phenomena seem to be independent of the details of the physical system at hand. Instead, only the system’s general features seem to matter [1142], such as (i) the spatial dimensionality (e.g., the system’s arrangement in a two-dimensional, three-dimensional, or more-dimensional lattice), (ii) the dimensionality of the order parameter (e.g., the system’s spin dimensionality), and (iii) the range of microscopic interactions (e.g., only first neighbours interact). Many physical quantities that describe the system’s state near a critical point have a power-law form whose main feature is the *critical exponent*. The ubiquity of power laws is often referred to as *universality*—different systems with the same values of the critical exponent are said to belong to the same universality class. In what follows, we focus on some of the critical phenomena specific to the Earth climate system.

Critical phenomena in atmospheric precipitation. Earth’s atmosphere is a fluid in complex motion that dynamically varies in space and time. Despite its dynamic complexity, from a meteorological perspective, the atmosphere is driven by slow large-scale forcing (moisture convergence, evaporation, and radiative cooling) and rapid convective-buoyancy release (small-scale convection). Convection intensifies above a critical point in the water vapour, causing the onset of heavy precipitation. This intensification is reflected in the average precipitation rate as a function of the water vapour, which exhibits a relatively simple power-law behaviour as predicted by the theory of critical phenomena [1143].

Using satellite data from the Tropical Rainfall Measuring Mission, Ref. [1143] analysed the relationship between the precipitation rate, P , and the water vapour, w . Various major ocean basins were covered by oceanic grid points between 20°S – 20°N . The precipitation and water-vapour data were collected at 0.25° latitude-longitude resolution. From a statistical physics perspective, quantities P and w were regarded as the *order parameter* and *tuning parameter*, respectively. It was found that, when the tuning parameter crosses its critical value, w_c , the order parameter can be well approximated by a power-law of the form

$$\langle P \rangle(w) = a(w - w_c)^\beta, \quad (143)$$

where a is a system-dependent constant and β is a critical exponent. The operator $\langle \cdot \rangle$ refers to averaging over all observations in a given region. The same power-law fits the data irrespective of the climatic region (Fig. 77A).

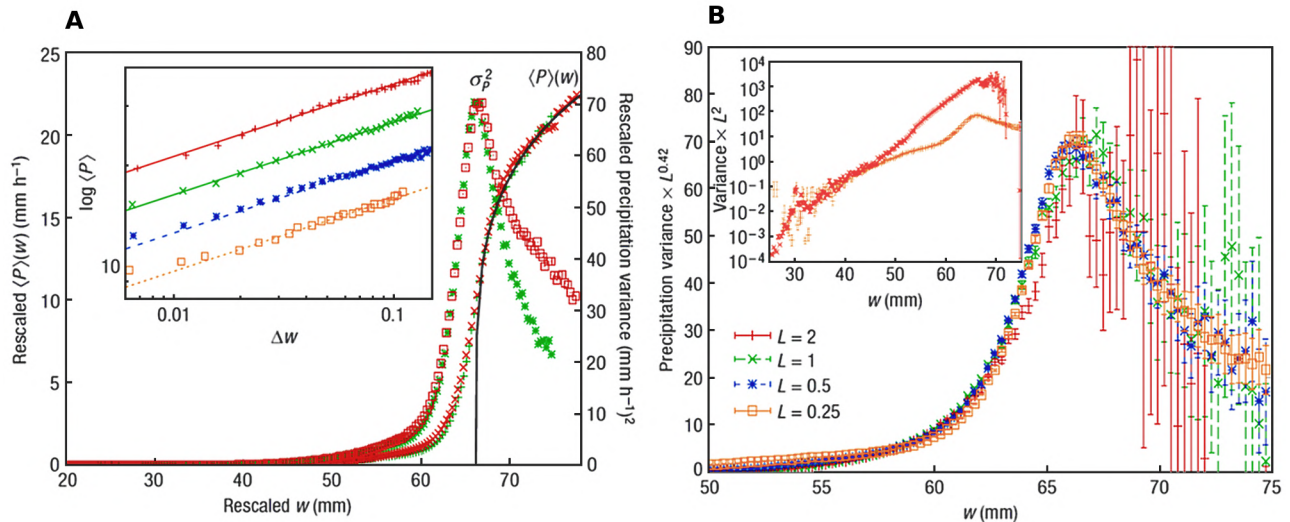


Figure 77: Critical phenomena in atmospheric precipitation. **A**, The average precipitation rates $\langle P \rangle(w)$ (order parameter) and their variances $\sigma_p^2(w)$ (susceptibility) are shown as a function of the water vapour (tuning parameter) for the eastern (red, 170°W – 70°W) and the western (green, 120°E – 170°W) Pacific Ocean. The solid curve stands for a power-law fit above the critical point. The inset shows, using double-logarithmic scales, $\langle P \rangle(w)$ as a function of the reduced water vapour, $\Delta w \equiv (w - w_c)/w_c$, for the western Pacific (green, 120°E – 170°W), the eastern Pacific (red, 170°W – 70°W), the Atlantic (blue, 70°W – 20°E), and the Indian Ocean (pink, 30°E – 120°E). Note the same slope irrespective of the climatic region. **B**, Finite-size scaling of the variance $\sigma_p^2(w; L)$ of the order parameter in the western Pacific. Near the critical point, $w > 57$ mm, the collapse of the curves is good, indicating $\sigma_p^2(w_c; L) \propto L^{-0.42}$, as expected from the theory of critical phenomena. The inset shows that relatively far away from the critical point, $w < 40$ mm, trivial scaling $\sigma_p^2(w; L) \propto L^{-2}$ works adequately. The error bars represent standard errors. Source: Reprinted figure from Ref. [1143].

Similarly, the critical exponent is universal and independent of the climatic region, with the value of 0.215 ± 0.02 (inset in Fig. 77A).

The susceptibility of the system, $\chi(w; L)$, was defined by means of the variance of the order parameter such that

$$\chi(w; L) = L^d \sigma_p^2(w; L), \quad (144)$$

where d stands for the system's dimensionality and L for the spatial resolution. Near the critical point w_c , however, the theory of critical phenomena suggests that [1144]

$$\chi(w; L) = L^{\gamma/\nu} \tilde{\chi}(\Delta w L^{1/\nu}), \quad (145)$$

where γ and ν are the standard critical exponents, $\Delta w \equiv (w - w_c)/w_c$ is the reduced water vapour, and $\tilde{\chi}(x)$ is the usual finite-size scaling function. When $\Delta w = 0$ (i.e., $w = w_c$), $\tilde{\chi}(0)$ is constant, implying the scaling relationship $\sigma_P^2(w_c; L) \propto L^{\gamma/\nu-d}$. Data confirm this relationship with the critical exponent ratio of $\gamma/\nu = 1.58$ (Fig. 77B). For $w > 57$ mm, all data indeed collapse into a single function. Furthermore, relatively far from the critical point, $w < 40$ mm, scaling is expectedly trivial (inset in Fig. 77B). These results show that the balance between slow large-scale forcing via moisture convergence, evaporation, and radiative cooling and rapid convective buoyancy release via small-scale convection leads to continuous (i.e., second order) phase transition such that below the critical point, there is very little precipitation, but once the critical point is crossed, precipitation rapidly increases with the water vapour. Interestingly, the balance between forcing and buoyancy release is stable, further suggesting that the described atmospheric criticality is in fact self-organised.

Hadley cell and percolation. The Hadley cell is a global-scale three-dimensional tropical atmospheric circulation that transports energy and angular momentum poleward. This circulation enables, among others, the trade winds, hurricanes, and the jet streams. The locations of the subtropical dry zones and the major tropical and subtropical deserts are strongly associated with the subsiding branches of the Hadley cell [1145]. Therefore, understanding how structure and intensity of the Hadley cell may change under global warming has attracted widespread attention. For example, an analysis of satellite observations indicated a poleward expansion, by $\approx 2^\circ$, of the Hadley cell over the period from 1979 to 2005 [1146]. A physical mechanism for the expansion of the Hadley cell was proposed shortly afterwards [1147], followed by a discovery of a robust weakening of the Hadley cell in the 21st century through the analysis of 30 different CMIP5 coupled model simulations [1148]. Observations, theory, and climate models are thus coming together to suggest the poleward expansion and weakening of the Hadley cell under global warming.

A standard approach to determining the strength of the Hadley cell is to compute the observed zonal-mean mass-stream function, Ψ . This function relates to the zonal-mean meridional wind velocity V via

$$[\bar{V}] = \frac{g}{2\pi R \cos \phi} \frac{\partial \Psi}{\partial p}, \quad (146)$$

where the operators $\bar{\cdot}$ and $[\cdot]$ stand for temporal and zonal averaging, respectively. The quantity g is the gravitational acceleration, R is the mean Earth

radius, ϕ is the latitude, and p designates pressure coordinates. When calculating the Ψ field, it is common to assume $\Psi = 0$ at the top of the atmosphere. Based on Eq. (146), the edges of the Hadley cell are identified as the first latitude poleward of the maximum of Ψ_{500} at which $\Psi_{500} = 0$, where the index 500 indicates the value of the stream function at 500 hPa [1147]. Although this conventional analysis has been applied to investigate the structure and intensity of the Hadley cell, there are some important limitations: (i) the latitude-longitude structure of the Hadley cell is not fully resolved because Eq. (146) only accounts for the latitudinal direction, and (ii) in contrast to theory and models that predict the decreasing intensity of the Hadley cell, the reanalysis datasets point to an increasing intensity [1149].

To circumvent the limitations of the conventional approach, Ref. [1122] analysed the structure and intensity of the Hadley cell using climate networks and *percolation* theory. The main question of percolation theory [1150] can be posed in several different ways, but in the context of network science, one seeks the probability q of node failure such that after $q\%$ of nodes do fail, the network changes from being connected to being disconnected. It turns out that there exist a critical probability q_c below which (i.e., for $q < q_c$) the network remains connected with probability one, but above which (i.e., for $q > q_c$) the network gets disconnected with probability one. This criticality strictly holds only for infinite networks, but is, in fact, easily observed in practice in networks with $\mathcal{O}(100)$ nodes. In Ref. [1122], the near-surface atmosphere was represented with a two-dimensional grid of points that turned into a lattice by adding links between nearest neighbours. These links were added as follows. First, the strength of each link, $W_{i,j}$, was calculated based on Eq. (134). Link strengths were then sorted in descending order. The strongest link was the first one to be added, then the second strongest, the third strongest, and so on. The resulting lattice-shaped climate network was found to undergo an abrupt phase transition in the order parameter G_1 , defined as the largest connected network component. Because the original grid points were embedded into the spherical Earth surface, the right expression for the order parameter was

$$G_1(M) = \frac{\max \left[\sum_{i \in S_1(M)} \cos(\phi_i), \dots, \sum_{i \in S_m(M)} \cos(\phi_i), \dots \right]}{\sum_{i=1}^N \cos(\phi_i)}, \quad (147)$$

where M is the number of added links, ϕ_i is the latitude of grid point i , and S_j is the j th connected network component in terms of the number of nodes. The percolation threshold at $M = M_c$ was determined by recording jumps in $G_1(M)$ with each added link, and singling out the largest jump to mark the value $G_c = G_1(M_c)$ and the corresponding critical link weight W_c [1151]. By altering the resolution of grid points, the theory of critical phenomena was used to confirm that the order parameter is indeed discontinuous at the percolation threshold and thus consistent with a first-order phase transition. The largest connected component of the climate network at the percolation threshold, obtained upon applying the described methodology, is located in the tropics, as expected from an analogue of the Hadley cell.

The purpose here was more than just finding a climate-network analogue of the Hadley cell. To determine the temporal evolution of the quantities G_c and W_c , a sequence of climate networks was constructed using successive and non-overlapping temporal windows with the length of 60 mos. The results could be fitted adequately with simple linear relationships

$$G_c(t) = a + \xi_G t, \quad (148a)$$

$$W_c(t) = b + \xi_W t, \quad (148b)$$

where a and b are constants, while ξ_G and ξ_W are the rates of change of the quantities G_c and W_c . Denoting the analogous rates of change for the Hadley cell, obtained via the conventional approach, with ξ_{ϕ_H} and ξ_{Ψ} , the results show a consistent expansion from the tropics poleward of the largest connected network component at the percolation threshold, as well as the weakening of the corresponding critical link weight. The same holds for the 31 CMIP5 21st century climate models and the reanalysis data (ERA-Interim and ERA-40). Put more quantitatively, most of the CMIP5 models exhibit $\xi_G > 0$ and $\xi_W < 0$, and similarly $\xi_{\phi_H} > 0$ and $\xi_{\Psi} < 0$ irrespective of the the climate scenario (Fig. 78, cf. panels A–C and D–F). The results obtained via climate-network analysis and using the conventional approach are highly correlated (Fig. 78, panels G–I).

The poleward expansion of the Hadley cell may result in (i) a drier future in some tropical or subtropical regions [1147] and (ii) a poleward migration of the location of the maximum tropical-cyclone intensity [1152]. The climate-network analysis described herein may therefore help to identify regions that are more probable to experience precipitation decline or hurricane intensification. Among the prime candidate regions to be affected by the Hadley-cell

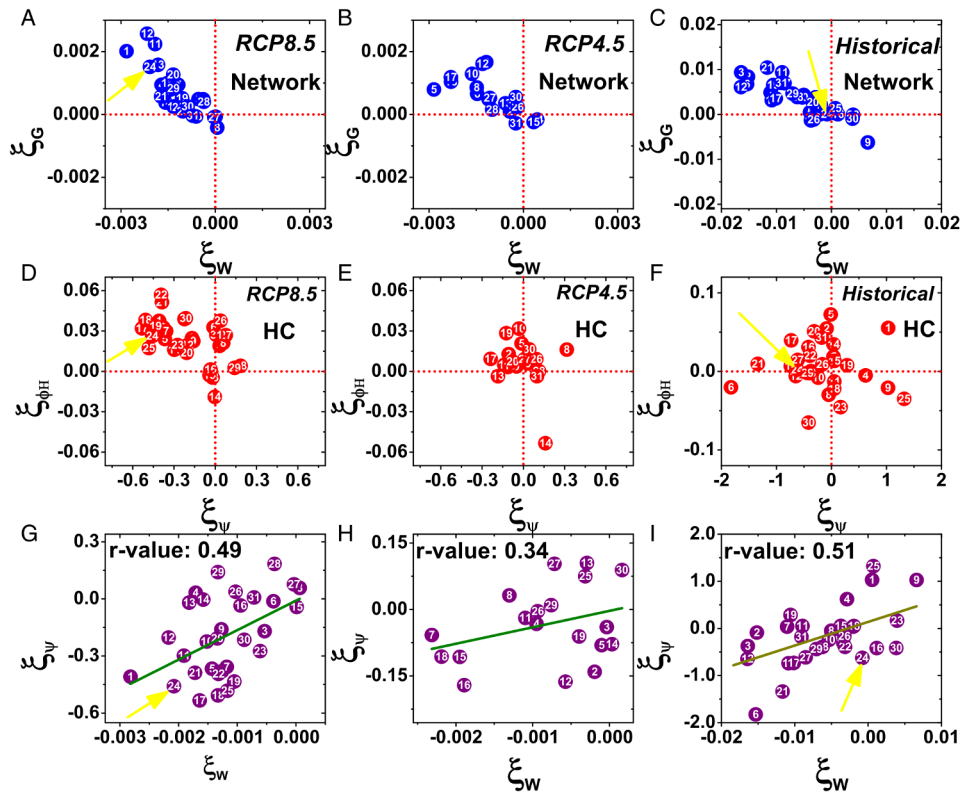


Figure 78: Comparing the rates of change of the Hadley cell and its climate-network analogue under three climate-change scenarios. The climate network analogue of the Hadley cell is the largest connected network component at the percolation threshold. Panels A–C show that this component is getting larger over time, while the critical link weight is decreasing under all scenarios for most of the CMIP5 climate models. Panels D–F show that the size of the Hadley cell and its intensity exhibit qualitatively similar behaviour as their climate-network analogues, although there is somewhat more ambiguity in the results, especially in the Historical scenario. Panels G–I show that the rates of change obtained via the climate-network analysis and the conventional approach show significant correlation. Numbering in the circles indexes the 31 CMIP5 climate models. Source: Reprinted figure from Ref. [1122].

expansion are northern India, southern Africa, and western Australia. Local governments in these and other potentially exposed regions should keep a close eye on climate science and take risk-mitigating actions until there is still time.

12.5. Future outlook

This chapter started by outlining the social consequences of global climate change and briefly venturing into the history of climate modelling. Thereafter, two research topics central to the current climate-change discourse—climate networks and critical phenomena—were introduced. Both of these topics originate from statistical physics, thus sharing their origins with many other themes in the present review that fall more squarely into the domain of social physics. The reason for the shared origins of the climate-change research and the social-physics research in the stricter sense is that the climate system is an epitome of complexity as much as human society is. Despite the fast-paced progress seen over the past two decades, there is still a lot of work ahead, especially in the context of integrating climate predictions into social dynamics.

Because climate networks are constructed by applying similarity measures to observational data, the underlying physical mechanisms and processes often remain hidden or unclear. Shedding light on such mechanisms and processes may, however, substantially impact our understanding of climate change and subsequently improve the predictive power of numerical climate models. A promising methodology that has emerged in recent years and could play an instrumental role in demystifying climate change is machine learning and AI [1153]. Considering the inherent ‘black box’ structure of climate systems, the integration of AI and visual analytics provides a potential solution [1154].

Another key issue is the question of analysing the climate resilience of ecosystems and economies. Currently, for example, there is a lack of appropriate models to fully understand and predict the effects of cascading failures [266], triggered by extreme climate and weather events, on critical interdependent infrastructures. Closing this knowledge gap is a crucial step towards climate-resilient society.

13. Epilogue: Keeping the dialogue open

We hope this review has given the reader an overview of physicists’ contributions to multidisciplinary social science. To make the story contiguous, we had to sacrifice some topics, like the physics of art (music [1155], painting [1156], dance [1157], etc.), agriculture [1158], gastronomy [1159], ethnology (how ethnic groups remember their shared history) [1160], civil

unrest [1161], etc. Conversely, we imagine some readers finding our definition of social physics too generous, especially at the border between physics and engineering, artificial intelligence, and climate modelling. Perhaps a better title would be *Human physics*—‘human’ as in ‘topics affecting humans’.

We started our exposé arguing that physics has played a fundamental role in the modern movement towards multidisciplinary. However, physicists entering multidisciplinary research have a bad reputation for their imperious attitude: “Step aside! We’ll show you how it’s done.” When this happens, even if unintentional, it threatens mutual respect and understanding between collaborators and jeopardises the overall success of collaborative interactions. But what better way to ensure mutual respect and understanding than to keep the dialogue open.

To illustrate what we have in mind, a physicist’s strength lies in putting quantitative methods to good use, be it rigorous data analyses or complex numerical simulations. The use of quantitative methods, however, is preceded by formulating research hypotheses of interest or model assumptions of relevance to the problem at hand. Seeking inputs from experts is absolutely crucial in this stage because intuition and common sense cannot replace expert knowledge, and may easily lead to simplistic and naive hypotheses or assumptions. Accordingly, before quantitative methods are employed, physicists for the most part need to be on the receiving end of the dialogue with their multidisciplinary collaborators.

Another strength that is rather unique to physicists is seeing the big picture and consequently making approximations that simplify the problem, but still account for the main processes at play. It is important to recognise that such approximations and subsequent simplifications go against the training received by researchers from many other disciplines. Among ecologists, for example, the focus on biodiversity is so prevalent that general patterns often come secondary to exceptions. Accordingly, when quantitative methods are employed, physicists for the most part need to be on the transmitting end of the dialogue with their multidisciplinary collaborators.

Many more situations are bound to arise in practice in which keeping the dialogue open will be crucial to success. They demand patience and care, but when resolved satisfactorily, they lead to insightful and impactful research that is so much needed to ensure the continued prosperity of humankind.

Acknowledgements

We are thankful to Sheikh Taslim Ali, Yosef Ashkenazy, Simon Cauchemez, Xiaosong Chen, Benjamin J. Cowling, Zhanwei Du, Shlomo Havlin, Juer-gen Kurths, Josef Ludescher, Jun Meng, Lauren A. Meyers, Henrik Salje, Hans Joachim Schellnhuber, Maximilian Schich, Joseph T. Wu, and Xi-aoke Xu. M. J. was supported by the Japan Society for the Promotion of Science (JSPS) KAKENHI grants-in-aid for scientific research JP20H04288, JP21K04545, and JP21H03625. P. H. was supported by JSPS KAKENHI JP21H04595. K. K. was supported by the Japan Science and Technology Agency (JST), PRESTO grant JPMJPR20M2. Z. W. was supported by the National Natural Science Foundation for Distinguished Young Scholars (grant no. 62025602), the National Natural Science Foundation of China (grants nos. U1803263, 11931015, 81961138010), the Fok Ying-Tong Edu-cation Foundation of China (grant no. 171105), and the Key Technology Research-and-Development Program of the Science and Technology-Scientific and Technological Innovation Team of Shaanxi Province (grant no. 2020TD-013). S. G. and T. K. were supported by the Croatian Science Foundation project 2018-01-3150 AqADAPT. B. P. was supported by the Slovenian Re-search Agency (grant no. J7-3156). L. W. was supported by the European Research Council (grant no. 804744), the Engineering and Physical Sciences Research Council (EPSRC) Impact Acceleration Grant RG90413, Health and Medical Research Fund, Food and Health Bureau, Government of the Hong Kong Special Administrative Region (grant no. 20190712), National Natu-ral Science Foundation of China (grant no. 62173065 and 11975025), and the International COVID-19 Data Alliance (ICODA), an initiative funded by the COVID-19 Therapeutics Accelerator and convened by Health Data Research UK. W. L. was supported by The National University of Singapore Start-up and Ministry of Education Tier 1 grants (grant no. WBS R-109-000-270-133). J. F. was supported by the East Africa Peru India Climate Capacities—EPICC project, which is a part of the International Climate Ini-tiative (IKI) sponsored by the Federal Ministry for the Environment, Nature Conservation, and Nuclear Safety (BMU) on the basis of a decision adopted by the German Bundestag. The Potsdam Institute for Climate Impact Re-search (PIK) is leading the execution of the project together with project partners The Energy and Resources Institute (TERI) and the German Mete-orological Service, Deutscher Wetterdienst (DWD). M. P. was supported by the Slovenian Research Agency (grant nos. P1-0403, J1-2457, and J1-9112).

References

- [1] P. Ball, The physical modelling of society: a historical perspective, *Physica A* 314 (2002) 1–14.
- [2] R. Sinatra, P. Deville, M. Szell, D. Wang, A.-L. Barabási, A century of physics, *Nat. Phys.* 11 (2015) 791–796.
- [3] D. A. Bromley, *A century of physics*, Springer Science & Business Media, 2002.
- [4] G. Gabrielse, D. Hanneke, T. Kinoshita, M. Nio, B. Odom, New determination of the fine structure constant from the electron g value and QED, *Phys. Rev. Lett.* 97 (2006) 030802.
- [5] C. W. F. Everitt, D. B. DeBra, B. W. Parkinson, J. P. Turneaure, J. W. Conklin, M. I. Heifetz, G. M. Keiser, A. S. Silbergleit, T. Holmes, J. Kolodziejczak, et al., Gravity probe B: final results of a space experiment to test general relativity, *Phys. Rev. Lett.* 106 (2011) 221101.
- [6] R. M. Nisbet, E. B. Muller, K. Lika, S. A. L. M. Kooijman, From molecules to ecosystems through dynamic energy budget models, *J. Anim. Ecol.* 69 (2000) 913–926.
- [7] J. H. Brown, J. F. Gillooly, A. P. Allen, V. M. Savage, G. B. West, Toward a metabolic theory of ecology, *Ecology* 85 (2004) 1771–1789.
- [8] S. A. L. M. Kooijman, *Dynamic energy budget theory for metabolic organisation*, Cambridge University Press, 2010.
- [9] M. Jusup, T. Sousa, T. Domingos, V. Labinac, N. Marn, Z. Wang, T. Klanjšček, Physics of metabolic organization, *Phys. Life Rev.* 20 (2017) 1–39.
- [10] M. Kearney, W. Porter, Mechanistic niche modelling: combining physiological and spatial data to predict species’ ranges, *Ecol. Lett.* 12 (2009) 334–350.
- [11] M. R. Kearney, W. P. Porter, NicheMapR—an R package for biophysical modelling: the ectotherm and Dynamic Energy Budget models, *Ecography* 43 (2020) 85–96.

- [12] N. C. Grassly, C. Fraser, Mathematical models of infectious disease transmission, *Nat. Rev. Microbiol.* 6 (2008) 477–487.
- [13] G.-Q. Sun, M. Jusup, Z. Jin, Y. Wang, Z. Wang, Pattern transitions in spatial epidemics: Mechanisms and emergent properties, *Phys. Life Rev.* 19 (2016) 43–73.
- [14] D. Wodarz, M. A. Nowak, Mathematical models of HIV pathogenesis and treatment, *BioEssays* 24 (2002) 1178–1187.
- [15] C. A. Beauchemin, A. Handel, A review of mathematical models of influenza A infections within a host or cell culture: lessons learned and challenges ahead, *BMC Public Health* 11 (2011) 1–15.
- [16] D. M. Kaplan, C. F. Craver, The explanatory force of dynamical and mathematical models in neuroscience: A mechanistic perspective, *Philos. Sci.* 78 (2011) 601–627.
- [17] D. S. Bassett, P. Zurn, J. I. Gold, On the nature and use of models in network neuroscience, *Nat. Rev. Neurosci.* 19 (2018) 566–578.
- [18] F. Michor, Y. Iwasa, M. A. Nowak, Dynamics of cancer progression, *Nat. Rev. Cancer* 4 (2004) 197–205.
- [19] R. Eftimie, J. J. Gillard, D. A. Cantrell, Mathematical models for immunology: current state of the art and future research directions, *Bull. Math. Biol.* 78 (2016) 2091–2134.
- [20] R. P. Abelson, Mathematical models in social psychology, in: L. Berkowitz (Ed.), *Advances in experimental social psychology*, volume 3, Elsevier, 1967, pp. 1–54.
- [21] J. L. Rodgers, The epistemology of mathematical and statistical modeling: a quiet methodological revolution., *Am. Psychol.* 65 (2010) 1–12.
- [22] A. B. Sørensen, Mathematical models in sociology, *Annu. Rev. Sociol.* 4 (1978) 345–371.
- [23] P. J. Carrington, J. Scott, S. Wasserman, *Models and methods in social network analysis*, Cambridge University Press, 2005.

- [24] S. S. Schweber, Physics, community and the crisis in physical theory, *Phys. Today* 46 (1993) 34–34.
- [25] J. Lykken, M. Spiropulu, Supersymmetry and the crisis in physics, *Sci. Am.* 310 (2014) 34–39.
- [26] L. Smolin, *The trouble with physics: the rise of string theory, the fall of a science, and what comes next*, Houghton Mifflin Harcourt, 2007.
- [27] S. Hossenfelder, *Screams for explanation: finetuning and naturalness in the foundations of physics*, 2019. e-print arXiv:1801.02176.
- [28] D. Alves, N. Arkani-Hamed, S. Arora, Y. Bai, M. Baumgart, J. Berger, M. Buckley, B. Butler, S. Chang, H.-C. Cheng, et al., Simplified models for LHC new physics searches, *J. Phys. G* 39 (2012) 105005.
- [29] M. Safronova, D. Budker, D. DeMille, D. F. J. Kimball, A. Derevianko, C. W. Clark, Search for new physics with atoms and molecules, *Rev. Mod. Phys.* 90 (2018) 025008.
- [30] J. D. Farmer, Physicists attempt to scale the ivory towers of finance, *Int. J. Theor. Appl. Finance* 3 (2000) 311–333.
- [31] United Nations, Department of Economic and Social Affairs, Population Division, *World Urbanization Prospects: The 2018 Revision*, 2018. Online edition. Available at: <https://population.un.org/wup/Download/>.
- [32] R. W. Caves, *Encyclopedia of the City*, Taylor & Francis, 2005.
- [33] M. P. Cohen, *Cities in times of crisis. The response of local governments in light of the global economic crisis: the role of the formation of human capital, urban innovation and strategic planning*, 2011. Institute of Urban and Regional Development (IURD) Working Paper No. 2011-01. Available at: <https://escholarship.org/uc/item/3432p4rb>.
- [34] M. Perc, J. J. Jordan, D. G. Rand, Z. Wang, S. Boccaletti, A. Szolnoki, Statistical physics of human cooperation, *Phys. Rep.* 687 (2017) 1–51.
- [35] U. Harush, B. Barzel, Dynamic patterns of information flow in complex networks, *Nat. Commun.* 8 (2017) 2181.

- [36] J. Van Dijk, Mafia markers: assessing organized crime and its impact upon societies, *Trends Organ. Crime* 10 (2007) 39–56.
- [37] R. B. Taylor, The impact of crime on communities, *Ann. Am. Acad. Political Soc. Sci.* 539 (1995) 28–45.
- [38] J. Bijak, D. Kupiszewska, M. Kupiszewski, Replacement migration revisited: Simulations of the effects of selected population and labor market strategies for the aging Europe, 2002–2052, *Popul. Res. Policy Rev.* 27 (2008) 321–342.
- [39] J. G. Williamson, Demographic dividends revisited, *Asian Dev. Rev.* 30 (2013) 1–25.
- [40] C. Postelnicescu, Europe’s new identity: The refugee crisis and the rise of nationalism, *Eur. J. Psychol* 12 (2016) 203–209.
- [41] S. M. Holmes, H. Castañeda, Representing the “European refugee crisis” in Germany and beyond: Deservingness and difference, life and death, *Am. Ethnol.* 43 (2016) 12–24.
- [42] H. V. Fineberg, Pandemic preparedness and response—lessons from the H1N1 influenza of 2009, *N. Engl. J. Med.* 370 (2014) 1335–1342.
- [43] R. I. L. Eggen, R. Behra, P. Burkhardt-Holm, B. I. Escher, N. Schweigert, Challenges in ecotoxicology, *Environ. Sci. Technol.* 38 (2004) 58A–64A.
- [44] A. P. Kirilenko, R. A. Sedjo, Climate change impacts on forestry, *Proc. Natl. Acad. Sci. USA* 104 (2007) 19697–19702.
- [45] M. E. S. Louis, J. J. Hess, Climate change: impacts on and implications for global health, *Am. J. Prev. Med.* 35 (2008) 527–538.
- [46] T. Wheeler, J. Von Braun, Climate change impacts on global food security, *Science* 341 (2013) 508–513.
- [47] M. Batty, The size, scale, and shape of cities, *Science* 319 (2008) 769–771.
- [48] M. Batty, *The new science of cities*, The MIT Press, 2013.

- [49] M. Barthelemy, The statistical physics of cities, *Nat. Rev. Phys.* 1 (2019) 406–415.
- [50] E. Arcaute, E. Hatna, P. Ferguson, H. Youn, A. Johansson, M. Batty, Constructing cities, deconstructing scaling laws, *J. R. Soc. Interface* 12 (2015) 20140745.
- [51] H. D. Rozenfeld, D. Rybski, J. S. Andrade, M. Batty, H. E. Stanley, H. A. Makse, Laws of population growth, *Proc. Natl. Acad. Sci. USA* 105 (2008) 18702–18707.
- [52] G. K. Zipf, *Human behavior and the principle of least effort*, Addison-Wesley Press, 1949.
- [53] D. H. Zanette, S. C. Manrubia, Role of intermittency in urban development: a model of large-scale city formation, *Phys. Rev. Lett.* 79 (1997) 523.
- [54] K. T. Soo, Zipf’s law for cities: a cross-country investigation, *Reg. Sci. Urban. Econ.* 35 (2005) 239–263.
- [55] Y. Malevergne, V. Pisarenko, D. Sornette, Testing the pareto against the lognormal distributions with the uniformly most powerful unbiased test applied to the distribution of cities, *Phys. Rev. E* 83 (2011) 036111.
- [56] M. Marsili, Y.-C. Zhang, Interacting individuals leading to Zipf’s law, *Phys. Rev. Lett.* 80 (1998) 2741.
- [57] H. A. Simon, On a class of skew distribution functions, *Biometrika* 42 (1955) 425–440.
- [58] R. K. Merton, The Matthew effect in science: The reward and communication systems of science are considered, *Science* 159 (1968) 56–63.
- [59] D. S. Price, A general theory of bibliometric and other cumulative advantage processes, *J. Am. Soc. Inf. Sci.* 27 (1976) 292–306.
- [60] X. Gabaix, Zipf’s law for cities: an explanation, *Q. J. Econ.* 114 (1999) 739–767.

- [61] M. E. J. Newman, Power laws, Pareto distributions and Zipf’s law, *Contemp. Phys.* 46 (2005) 323–351.
- [62] E. H. Decker, A. J. Kerkhoff, M. E. Moses, Global patterns of city size distributions and their fundamental drivers, *PLOS ONE* 2 (2007) e934.
- [63] S. C. Manrubia, D. H. Zanette, R. V. Sole, Transient dynamics and scaling phenomena in urban growth, *Fractals* 7 (1999) 1–8.
- [64] T. A. Witten, L. M. Sander, Diffusion-limited aggregation, a kinetic critical phenomenon, *Phys. Rev. Lett.* 47 (1981) 1400–1403.
- [65] C. Andersson, K. Lindgren, S. Rasmussen, R. White, Urban growth simulation from “first principles”, *Phys. Rev. E* 66 (2002) 026204.
- [66] H. A. Makse, S. Havlin, H. E. Stanley, Modelling urban growth patterns, *Nature* 377 (1995) 608–612.
- [67] K. C. Seto, B. Güneralp, L. R. Hutyrá, Global forecasts of urban expansion to 2030 and direct impacts on biodiversity and carbon pools, *Proc. Natl. Acad. Sci. USA* 109 (2012) 16083–16088.
- [68] C. A. Doxiadis, *Ekistics and traffic*, *Traffic Q.* 17 (1963) 439–457.
- [69] R. Abler, J. S. Adams, P. Gould, *Spatial organization: the geographer’s view of the world*, Prentice-Hall, 1971.
- [70] W. Christaller, *Central Places in Southern Germany*, Prentice-Hall, 1966.
- [71] A. Losch, *The Economics of Location*, Yale University Press, 1954.
- [72] J. B. Parr, Structure and size in the urban system of Lösch, *Econ. Geogr.* 49 (1973) 185–212.
- [73] M. Barthélemy, A. Flammini, Modeling urban street patterns, *Phys. Rev. Lett.* 100 (2008) 138702.
- [74] M. Barthélemy, *The Structure and Dynamics of Cities: Urban Data Analysis and Theoretical Modeling*, Cambridge University Press, 2016.

- [75] L. M. A. Bettencourt, *Introduction to Urban Science: Evidence and Theory of Cities as Complex Systems*, MIT Press, Cambridge MA, 2021.
- [76] M. Barthélemy, *Morphogenesis of spatial networks*, Springer, 2018.
- [77] M. T. Gastner, M. E. J. Newman, Shape and efficiency in spatial distribution networks, *J. Stat. Mech.* 2006 (2006) P01015.
- [78] A. Fabrikant, E. Koutsoupias, C. H. Papadimitriou, Heuristically optimized trade-offs: A new paradigm for power laws in the Internet, in: P. Widmayer, S. Eidenbenz, F. Triguero, R. Morales, R. Conejo, M. Hennessy (Eds.), *International Colloquium on Automata, Languages, and Programming*, volume 2380 of *Lecture Notes in Computer Science*, Springer, 2002, pp. 110–122.
- [79] T. C. Schelling, Dynamic models of segregation, *J. Math. Sociol.* 1 (1971) 143–186.
- [80] D. Vinković, A. Kirman, A physical analogue of the Schelling model, *Proc. Natl. Acad. Sci. USA* 103 (2006) 19261–19265.
- [81] L. Dall’Asta, C. Castellano, M. Marsili, Statistical physics of the Schelling model of segregation, *J. Stat. Mech.* 2008 (2008) L07002.
- [82] D. Stauffer, A biased review of sociophysics, *J. Stat. Phys.* 151 (2013) 9–20.
- [83] P. Bak, *How Nature Works: The Science of Self-Organized Criticality*, Copernicus, 1996.
- [84] G. B. West, *Scale: The Universal Laws of Life and Death in Organisms, Cities and Companies*, Weidenfeld & Nicolson, 2017.
- [85] G. B. West, J. H. Brown, B. J. Enquist, A general model for the origin of allometric scaling laws in biology, *Science* 276 (1997) 122–126.
- [86] L. M. Bettencourt, J. Lobo, D. Helbing, C. Kühnert, G. B. West, Growth, innovation, scaling, and the pace of life in cities, *Proc. Natl. Acad. Sci. USA* 104 (2007) 7301–7306.

- [87] J. Um, S.-W. Son, S.-I. Lee, H. Jeong, B. J. Kim, Scaling laws between population and facility densities, *Proc. Natl. Acad. Sci. USA* 106 (2009) 14236–14240.
- [88] L. M. A. Bettencourt, The origins of scaling in cities, *Science* 340 (2013) 1438–1441.
- [89] E. G. Ravenstein, The laws of migration, *J. Stat. Soc. Lond.* 48 (1885) 167–235.
- [90] G. K. Zipf, The P1 P2 / D hypothesis: On the intercity movement of persons, *Am. Soc. Rev.* 11 (1946) 677–686.
- [91] Y. Chen, L. Huang, A scaling approach to evaluating the distance exponent of the urban gravity model, *Chaos Solitons Fractals* 109 (2018) 303–313.
- [92] M. Lee, P. Holme, Relating land use and human intra-city mobility, *PLOS ONE* 10 (2015) e0140152.
- [93] S. A. Stouffer, Intervening opportunities: A theory relating mobility and distance, *Am. Soc. Rev.* 5 (1940) 845–867.
- [94] M. C. Gonzalez, C. A. Hidalgo, A.-L. Barabási, Understanding individual human mobility patterns, *Nature* 453 (2008) 779–782.
- [95] J. Park, D.-S. Lee, M. C. González, The eigenmode analysis of human motion, *J. Stat. Mech.* 2010 (2010) P11021.
- [96] C. Song, Z. Qu, N. Blumm, A.-L. Barabási, Limits of predictability in human mobility, *Science* 327 (2010) 1018–1021.
- [97] X. Lu, L. Bengtsson, P. Holme, Predictability of population displacement after the 2010 Haiti earthquake, *Proc. Natl. Acad. Sci. USA* 109 (2012) 11576–11581.
- [98] R. Louf, P. Jensen, M. Barthélemy, Emergence of hierarchy in cost-driven growth of spatial networks, *Proc. Natl. Acad. Sci. USA* 110 (2013) 8824–8829.
- [99] M. Lee, S. Cheon, S. Lee, Imbalance of pairwise efficiency in urban street network, 2018. e-print arXiv:1808.05844.

- [100] M. Lee, H. Barbosa, H. Youn, P. Holme, G. Ghoshal, Morphology of travel routes and the organization of cities, *Nat. Commun.* 8 (2017) 2229.
- [101] M. de Berg, M. van Kreveld, M. Overmars, O. Schwarzkopf, *Computational Geometry: Algorithms and Applications*, Springer, 1997, pp. 305–315.
- [102] S. H. Lee, P. Holme, Exploring maps with greedy navigators, *Phys. Rev. Lett.* 108 (2012) 128701.
- [103] T. Wolbers, M. Hegarty, What determines our navigational abilities?, *Trends Cogn. Sci.* 14 (2010) 138–146.
- [104] W. Luo, Y. Wang, X. Liu, S. Gao, Cities as spatial and social networks: towards a spatio-socio-semantic analysis framework, in: X. Ye, X. Liu (Eds.), *Cities as Spatial and Social Networks*, Springer, 2019, pp. 21–37.
- [105] D. Ni, *Traffic Flow Theory: Characteristics, Experimental Methods, and Numerical Techniques*, Butterworth Heinemann, 2015.
- [106] F. Kessels, *Traffic Flow Modelling: Introduction to Traffic Flow Theory Through a Genealogy of Models*, Springer, 1999.
- [107] B. D. Greenshields, J. R. Bibbins, W. S. Channing, H. H. Miller, A study of traffic capacity, in: *Highway Research Board Proceedings 14*, Highway Research Board, 1935, pp. 448–477.
- [108] D. Chowdhury, L. Santen, A. Schadschneider, Statistical physics of vehicular traffic and some related systems, *Phys. Rep.* 329 (2000) 199–329.
- [109] D. Helbing, Traffic and related self-driven many-particle systems, *Rev. Mod. Phys.* 73 (2001) 1067–1141.
- [110] T. Nagatani, The physics of traffic jams, *Rep. Prog. Phys.* 65 (2002) 1331–1386.
- [111] A. S. Mihaita, H. Li, Z. He, M.-A. Rizoïu, Motorway traffic flow prediction using advanced deep learning, in: *Proceedings of the 2019*

IEEE Intelligent Transportation Systems Conference (ITSC), IEEE, 2019, pp. 1683–1690.

- [112] F. van Wageningen-Kessels, H. van Lint, K. Vuik, S. Hoogendoorn, Genealogy of traffic flow models, *EURO J. Transp. Logist.* 4 (2015) 445–473.
- [113] M. J. Lighthill, G. B. Whitham, On kinematic waves II. A theory of traffic flow on long crowded roads, *Proc. R. Soc. A* 229 (1955) 317–345.
- [114] A. Reuschel, Fahrzeugbewegungen in der Kolonne, *Österreichisches Ingenieur Archiv* 4 (1950) 193–215.
- [115] K. Nagel, M. Schreckenberg, A cellular automaton model for freeway traffic, *J. Phys. I* 2 (1992) 2221–2229.
- [116] F. L. Hall, B. L. Allen, M. A. Gunter, Empirical analysis of freeway flow-density relationships, *Transp. Res. A* 20 (1986) 197–210.
- [117] B. S. Kerner, *Breakdown in Traffic Networks: Fundamentals of Traffic Science*, Springer, 2017.
- [118] B. S. Kerner, H. Rehborn, Experimental properties of complexity in traffic flow, *Phys. Rev. E* 53 (1996) R4275–R4278.
- [119] S.-i. Tadaki, M. Kikuchi, M. Fukui, A. Nakayama, K. Nishinari, A. Shibata, Y. Sugiyama, T. Yosida, S. Yukawa, Phase transition in traffic jam experiment on a circuit, *New J. Phys.* 15 (2013) 103034.
- [120] J. Treiterer, J. A. Myers, The hysteresis phenomenon in traffic flow, in: D. Buckley (Ed.), *Proceedings of the 6th International Symposium on Transportation and Traffic Theory*, Reed, 1974, pp. 13–38.
- [121] W. Leutzbach, *Introduction to the Theory of Traffic Flow*, Springer, 1988.
- [122] Y. Sugiyama, M. Fukui, M. Kikuchi, K. Hasebe, A. Nakayama, K. Nishinari, S. Tadaki, S. Yukawa, Traffic jams without bottlenecks: Experimental evidence for the physical mechanism of the formation of a jam, *New J. Phys.* 10 (2008) 033001.

- [123] A. Kesting, M. Treiber, D. Helbing, Agents for traffic simulation, in: A. M. Uhrmacher, D. Weyns (Eds.), *Multi-agent systems: Simulation and applications*, CRC Press, 2018, pp. 325–356.
- [124] B. S. Kerner, Experimental features of self-organization in traffic flow, *Phys. Rev. Lett.* 81 (1998) 3797–3800.
- [125] I. Prigogine, F. C. Andrews, A Boltzmann-like approach for traffic flow, *Oper. Res.* 8 (1960) 789–797.
- [126] S. Paveri-Fontana, On Boltzmann-like treatments for traffic flow: a critical review of the basic model and an alternative proposal for dilute traffic analysis, *Transport. Res.* 9 (1975) 225–235.
- [127] S. Yukawa, M. Kikuchi, Coupled-map modeling of one-dimensional traffic flow, *J. Phys. Soc. Jpn.* 64 (1995) 35–38.
- [128] S. Krauß, P. Wagner, C. Gawron, Continuous limit of the Nagel-Schreckenberg model, *Phys. Rev. E* 54 (1996) 3707–3712.
- [129] T. Horiguchi, T. Sakakibara, Numerical simulations for traffic-flow models on a decorated square lattice, *Physica A* 252 (1998) 388–404.
- [130] D. Helbing, B. A. Huberman, Coherent moving states in highway traffic, *Nature* 396 (1998) 738–740.
- [131] M. Takayasu, H. Takayasu, $1/f$ noise in a traffic model, *Fractals* 1 (1993) 860–866.
- [132] C. T. MacDonald, J. H. Gibbs, A. C. Pipkin, Kinetics of biopolymerization on nucleic acid templates, *Biopolymers* 6 (1968) 1–25.
- [133] J. M. Burgers, A mathematical model illustrating the theory of turbulence, in: R. von Mises, T. von Karman (Eds.), *Advances in Applied Mechanics*, volume 1, Academic Press, 1948, pp. 171–199.
- [134] D. Forster, D. R. Nelson, M. J. Stephen, Long-time tails and the large-eddy behavior of a randomly stirred fluid, *Phys. Rev. Lett.* 36 (1976) 867.

- [135] O. Biham, A. A. Middleton, D. Levine, Self-organization and a dynamical transition in traffic-flow models, *Phys. Rev. A* 46 (1992) R6124–R6127.
- [136] D. Helbing, J. Keltsch, P. Molnar, Modelling the evolution of human trail systems, *Nature* 388 (1997) 47–50.
- [137] C. Feliciani, K. Nishinari, Empirical analysis of the lane formation process in bidirectional pedestrian flow, *Phys. Rev. E* 94 (2016) 032304.
- [138] D. Helbing, I. Farkas, T. Vicsek, Simulating dynamical features of escape panic, *Nature* 407 (2000) 487–490.
- [139] L. F. Henderson, The statistics of crowd fluids, *Nature* 229 (1971) 381–383.
- [140] R. Krajewski, J. Bock, L. Kloeker, L. Eckstein, The highD dataset: A drone dataset of naturalistic vehicle trajectories on German highways for validation of highly automated driving systems, in: 2018 21st International Conference on Intelligent Transportation Systems (ITSC), IEEE, 2018, pp. 2118–2125.
- [141] J. Halkias, J. Colyar, NGSIM—Interstate 80 Freeway Dataset, 2006. Official website: <https://www.fhwa.dot.gov/publications/research/operations/07030/index.cfm>. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [142] J. V. Andersen, A. Nowak, Financial markets as interacting individuals: Price formation from models of complexity, in: *An Introduction to Socio-Finance*, Springer, 2013, pp. 59–76.
- [143] P. Ormerod, Ten years after “Worrying trends in econophysics”: developments and current challenges, *Eur. Phys. J. Spec. Top.* 225 (2016) 3281–3291.
- [144] A. Dragulescu, V. M. Yakovenko, Statistical mechanics of money, *Eur. Phys. J. B* 17 (2000) 723–729.
- [145] H. Aoyama, Y. Fujiwara, Y. Ikeda, H. Iyetomi, W. Souma, *Econophysics and companies: statistical life and death in complex business networks*, Cambridge University Press, 2010.

- [146] D. Braha, B. Stacey, Y. Bar-Yam, Corporate competition: A self-organized network, *Soc. Netw.* 33 (2011) 219–230.
- [147] W. Luo, P. Yin, Q. Di, F. Hardisty, A. M. MacEachren, A geovisual analytic approach to understanding geo-social relationships in the international trade network, *PLOS ONE* 9 (2014) e88666.
- [148] D. Harmon, M. Lagi, M. A. de Aguiar, D. D. Chinellato, D. Braha, I. R. Epstein, Y. Bar-Yam, Anticipating economic market crises using measures of collective panic, *PLOS ONE* 10 (2015) e0131871.
- [149] F. Abergel, H. Aoyama, B. K. Chakrabarti, A. Chakraborti, N. Deo, D. Raina, I. Vodenska, *Econophysics and sociophysics: Recent progress and future directions*, Springer, 2017.
- [150] J. Wachs, J. Kertész, A network approach to cartel detection in public auction markets, *Sci. Rep.* 9 (2019) 10818.
- [151] R. N. Mantegna, Lévy walks and enhanced diffusion in Milan stock exchange, *Physica A* 179 (1991) 232–242.
- [152] B. B. Mandelbrot, *The fractal geometry of nature*, Times Books, 1982.
- [153] T. Vicsek, *Fractal growth phenomena*, World scientific, 1992.
- [154] H. Takayasu, *Fractals in the physical sciences*, Manchester University Press, 1990.
- [155] P. Bak, C. Tang, K. Wiesenfeld, Self-organized criticality: An explanation of $1/f$ noise, *Phys. Rev. Lett.* 59 (1987) 381.
- [156] H. Takayasu, H. Miura, T. Hirabayashi, K. Hamada, Statistical properties of deterministic threshold elements—the case of market price, *Physica A* 184 (1992) 127–134.
- [157] R. N. Mantegna, H. E. Stanley, Scaling behaviour in the dynamics of an economic index, *Nature* 376 (1995) 46–49.
- [158] H. Takayasu, A.-H. Sato, M. Takayasu, Stable infinite variance fluctuations in randomly amplified Langevin systems, *Phys. Rev. Lett.* 79 (1997) 966.

- [159] R. Mantegna, H. Stanley, *An Introduction to Econophysics: Correlations and Complexity in Finance*, Cambridge University Press, 1999.
- [160] D. Sornette, *Why stock markets crash: critical events in complex financial systems*, Princeton University Press, 2003.
- [161] J.-P. Bouchaud, M. Potters, *Theory of financial risk and derivative pricing: from statistical physics to risk management*, Cambridge University Press, 2003.
- [162] H. Takayasu, *Empirical science of financial fluctuations: The advent of econophysics*, Springer, 2013.
- [163] H. Takayasu, *The Application of Econophysics: Proceedings of the Second Nikkei Econophysics Symposium*, Springer, 2012.
- [164] M. Takayasu, T. Watanabe, H. Takayasu, *Econophysics approaches to large-scale business data and financial crisis*, Springer, 2010.
- [165] T. Hirabayashi, H. Takayasu, H. Miura, K. Hamada, The behavior of a threshold model of market price in stock exchange, *Fractals* 1 (1993) 29–40.
- [166] K. Yamada, H. Takayasu, M. Takayasu, Characterization of foreign exchange market using the threshold-dealer-model, *Physica A* 382 (2007) 340–346.
- [167] A.-H. Sato, H. Takayasu, Dynamic numerical models of stock market price: from microscopic determinism to macroscopic randomness, *Physica A* 250 (1998) 231–252.
- [168] R. F. Engle, Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation, *Econometrica* 50 (1982) 987–1007.
- [169] A.-H. Sato, H. Takayasu, Derivation of ARCH(1) process from market price changes based on deterministic microscopic multi-agent, in: H. Takayasu (Ed.), *Empirical Science of Financial Fluctuations*, Springer, 2002, pp. 171–178.
- [170] K. Yamada, H. Takayasu, T. Ito, M. Takayasu, Solvable stochastic dealer models for financial markets, *Phys. Rev. E* 79 (2009) 051120.

- [171] K. Matsunaga, K. Yamada, H. Takayasu, M. Takayasu, Construction of the spread dealer model and its application (in japanese), *Trans. Jpn. Soc. Artif. Intell.* 27 (2012) 365–375.
- [172] A. Ciacci, T. Sueshige, H. Takayasu, K. Christensen, M. Takayasu, The microscopic relationships between triangular arbitrage and cross-currency correlations in a simple agent based model of foreign exchange markets, *PLOS ONE* 15 (2020) e0234709.
- [173] Y. Aiba, N. Hatano, H. Takayasu, K. Marumo, T. Shimizu, Triangular arbitrage as an interaction among foreign exchange rates, *Physica A* 310 (2002) 467–479.
- [174] T. Ito, K. Yamada, M. Takayasu, H. Takayasu, Execution risk and arbitrage opportunities in the foreign exchange markets, 2020. National Bureau of Economic Research (NBER) Working Paper No. 26706. Available at: <https://www.nber.org/papers/w26706>.
- [175] M. Takayasu, H. Takayasu, M. P. Okazaki, Transaction interval analysis of high resolution foreign exchange data, in: H. Takayasu (Ed.), *Empirical Science of Financial Fluctuations*, Springer, 2002, pp. 18–25.
- [176] M. Takayasu, H. Takayasu, Self-modulation processes and resulting generic $1/f$ fluctuations, *Physica A* 324 (2003) 101–107.
- [177] M. Takayasu, T. Mizuno, T. Ohnishi, H. Takayasu, Temporal characteristics of moving average of foreign exchange markets, in: H. Takayasu (Ed.), *Practical Fruits of Econophysics*, Springer, 2006, pp. 29–32.
- [178] M. Takayasu, T. Mizuno, H. Takayasu, Potential force observed in market dynamics, *Physica A* 370 (2006) 91–97.
- [179] M. Takayasu, K. Watanabe, T. Mizuno, H. Takayasu, Theoretical base of the PUCK-model with application to foreign exchange markets, in: M. Takayasu, T. Watanabe, H. Takayasu (Eds.), *Econophysics Approaches to Large-Scale Business Data and Financial Crisis*, Springer, 2010, pp. 79–98.

- [180] M. Takayasu, H. Takayasu, Continuum limit and renormalization of market price dynamics based on PUCK model, *Prog. Theor. Phys. Supp.* 179 (2009) 1–7.
- [181] K. Watanabe, H. Takayasu, M. Takayasu, Observation of two types of behaviors of financial bubbles and the related higher-order potential forces, *Prog. Theor. Phys. Supp.* 179 (2009) 8–16.
- [182] K. Watanabe, H. Takayasu, M. Takayasu, Random walker in temporally deforming higher-order potential forces observed in a financial crisis, *Phys. Rev. E* 80 (2009) 056110.
- [183] S. Maslov, Simple model of a limit order-driven market, *Physica A* 278 (2000) 571–578.
- [184] J.-P. Bouchaud, M. Mézard, M. Potters, Statistical properties of stock order books: empirical results and models, *Quant. Finance* 2 (2002) 251–256.
- [185] Y. Yura, H. Takayasu, D. Sornette, M. Takayasu, Financial brownian particle in the layered order-book fluid and fluctuation-dissipation relations, *Phys. Rev. Lett.* 112 (2014) 098703.
- [186] Y. Yura, H. Takayasu, D. Sornette, M. Takayasu, Financial Knudsen number: Breakdown of continuous price dynamics and asymmetric buy-and-sell structures confirmed by high-precision order-book information, *Phys. Rev. E* 92 (2015) 042811.
- [187] M. G. Daniels, J. D. Farmer, L. Gillemot, G. Iori, E. Smith, Quantitative model of price diffusion and market friction based on trading as a mechanistic random process, *Phys. Rev. Lett.* 90 (2003) 108102.
- [188] E. Smith, J. D. Farmer, L. Gillemot, S. Krishnamurthy, Statistical theory of the continuous double auction, *Quant. Finance* 3 (2003) 481–514.
- [189] F. Slanina, *Essentials of econophysics modelling*, Oxford University Press, 2013.
- [190] J.-P. Bouchaud, J. Bonart, J. Donier, M. Gould, *Trades, quotes and prices: financial markets under the microscope*, Cambridge University Press, 2018.

- [191] K. Kanazawa, T. Sueshige, H. Takayasu, M. Takayasu, Derivation of the Boltzmann equation for financial Brownian motion: Direct observation of the collective motion of high-frequency traders, *Phys. Rev. Lett.* 120 (2018) 138301.
- [192] K. Kanazawa, T. Sueshige, H. Takayasu, M. Takayasu, Kinetic theory for financial Brownian motion from microscopic dynamics, *Phys. Rev. E* 98 (2018) 052317.
- [193] T. Sueshige, K. Kanazawa, H. Takayasu, M. Takayasu, Ecology of trading strategies in a forex market for limit and market orders, *PLOS ONE* 13 (2018) e0208332.
- [194] A. B. Schmidt, Ecology of the modern institutional spot FX: The EBS market in 2011, 2012. Social Science Research Network (SSRN) Paper No. 1984070. Available at: <https://ssrn.com/abstract=1984070>.
- [195] I. Poirer, High-frequency trading and the flash crash: structural weaknesses in the securities markets and proposed regulatory responses, *Hastings Bus. Law J.* 8 (2012) 445–472.
- [196] EBS dealing rules—appendix EBS market, 2020. Available at: <https://www.cmegroup.com/trading/market-tech-and-data-services/files/ebs-dealing-rules-appendix-ebs-market.pdf>. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [197] C. W. Gardiner, *Handbook of stochastic methods*, Springer, 1985.
- [198] H. Haken, *Synergetics: Introduction and advanced topics*, Springer Science & Business Media, 2013.
- [199] L. D. Landau, E. M. Lifshitz, *Course of theoretical physics volume 1: Mechanics*, Elsevier, 1976.
- [200] D. J. Evans, G. Morriss, *Statistical mechanics of nonequilibrium liquids*, Cambridge University Press, 2008.
- [201] J.-P. Hansen, I. R. McDonald, *Theory of simple liquids*, Elsevier, 1990.
- [202] N. G. Van Kampen, *Stochastic processes in physics and chemistry, volume 1*, Elsevier, 1992.

- [203] H. Spohn, Kinetic equations from hamiltonian dynamics: Markovian limits, *Rev. Mod. Phys.* 52 (1980) 569.
- [204] T. Lux, The stable Paretian hypothesis and the frequency of large returns: an examination of major German stocks, *Appl. Financ. Econ.* 6 (1996) 463–475.
- [205] F. M. Longin, The asymptotic distribution of extreme stock market returns, *J. Business* 69 (1996) 383–408.
- [206] D. M. Guillaume, M. M. Dacorogna, R. R. Davé, U. A. Müller, R. B. Olsen, O. V. Pictet, From the bird’s eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets, *Finance Stoch.* 1 (1997) 95–129.
- [207] V. Plerou, P. Gopikrishnan, L. A. N. Amaral, M. Meyer, H. E. Stanley, Scaling of the distribution of price fluctuations of individual companies, *Phys. Rev. E* 60 (1999) 6519.
- [208] T. Sueshige, D. Sornette, H. Takayasu, M. Takayasu, Classification of position management strategies at the order-book level and their influences on future market-price formation, *PLOS ONE* 14 (2019) e0220645.
- [209] J. D. Farmer, Market force, ecology and evolution, *Ind. Corp. Chang.* 11 (2002) 895–953.
- [210] J. D. Farmer, S. Skouras, An ecological perspective on the future of computer trading, *Quant. Finance* 13 (2013) 325–346.
- [211] T. Odean, Are investors reluctant to realize their losses?, *J. Finance* 53 (1998) 1775–1798.
- [212] M. Grinblatt, M. Keloharju, The investment behavior and performance of various investor types: a study of Finland’s unique data set, *J. Financ. Econ.* 55 (2000) 43–67.
- [213] Y.-T. Lee, Y.-J. Liu, R. Roll, A. Subrahmanyam, Order imbalances and market efficiency: Evidence from the Taiwan Stock Exchange, *J. Financ. Quant. Anal.* 39 (2004) 327–341.

- [214] D. M. de Lachapelle, D. Challet, Turnover, account value and diversification of real traders: evidence of collective portfolio optimizing behavior, *New J. Phys.* 12 (2010) 075039.
- [215] B. Toth, Z. Eisler, F. Lillo, J. Kockelkoren, J.-P. Bouchaud, J. D. Farmer, How does the market react to your order flow?, *Quant. Finance* 12 (2012) 1015–1024.
- [216] M. Tumminello, F. Lillo, J. Piilo, R. N. Mantegna, Identification of clusters of investors from their real trading activity in a financial market, *New J. Phys.* 14 (2012) 013041.
- [217] F. Lillo, S. Micciche, M. Tumminello, J. Piilo, R. N. Mantegna, How news affects the trading behaviour of different categories of investors in a financial market, *Quant. Finance* 15 (2015) 213–229.
- [218] F. Musciotto, L. Marotta, J. Piilo, R. N. Mantegna, Long-term ecology of investors in a financial market, *Palgrave Commun.* 4 (2018) 92.
- [219] S. A. Frank, All of life is social, *Curr. Biol.* 17 (2007) R648–R650.
- [220] J. M. Smith, *Evolution and the Theory of Games*, Cambridge University Press, 1982.
- [221] J. L. Sachs, U. G. Mueller, T. P. Wilcox, J. J. Bull, The evolution of cooperation, *Q. Rev. Biol.* 79 (2004) 135–160.
- [222] D. J. Cooper, J. H. Kagel, Other-regarding preferences, in: J. H. Kagel, A. E. Roth (Eds.), *The Handbook of Experimental Economics*, Volume 2, Princeton University Press, 2016, pp. 217–289.
- [223] J. Henrich, M. Muthukrishna, The origins and psychology of human cooperation, *Annual Review of Psychology* 72 (2021) 207–240.
- [224] E. Akçay, Deconstructing evolutionary game theory: coevolution of social behaviors with their evolutionary setting, *Am. Nat.* 195 (2020) 315–330.
- [225] R. M. Dawes, Social dilemmas, *Annu. Rev. Psychol.* 31 (1980) 169–193.

- [226] S. Kuhn, Prisoner's Dilemma, in: E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy (Winter 2019 Edition), Metaphysics Research Lab, Stanford University, 2019, pp. —. Available at: <https://plato.stanford.edu/archives/win2019/entries/prisoner-dilemma/>.
- [227] J. W. Weibull, Evolutionary game theory, MIT Press, 1997.
- [228] J. Tanimoto, Fundamentals of evolutionary game theory and its applications, Springer, 2015.
- [229] P. D. Taylor, L. B. Jonker, Evolutionary stable strategies and game dynamics, *Math. Biosci.* 40 (1978) 145–156.
- [230] M. Doebeli, C. Hauert, Models of cooperation based on the Prisoner's Dilemma and the Snowdrift game, *Ecol. Lett.* 8 (2005) 748–766.
- [231] Z. Wang, M. Jusup, S. Kokubo, J. Tanimoto, Universal scaling for the dilemma strength in evolutionary games, *Phys. Life Rev.* 14 (2015) 1–30.
- [232] H. Ito, J. Tanimoto, Scaling the phase-planes of social dilemma strengths shows game-class changes in the five rules governing the evolution of cooperation, *R. Soc. Open Sci.* 5 (2018) 181085.
- [233] M. R. Arefin, K. M. A. Kabir, M. Jusup, H. Ito, J. Tanimoto, Social efficiency deficit deciphers social dilemmas, *Sci. Rep.* 10 (2020) 16092.
- [234] M. A. Nowak, Five rules for the evolution of cooperation, *Science* 314 (2006) 1560–1563.
- [235] C. Taylor, M. A. Nowak, Transforming the dilemma, *Evolution* 61 (2007) 2281–2292.
- [236] I. Eshel, L. L. Cavalli-Sforza, Assortment of encounters and evolution of cooperativeness, *Proc. Natl. Acad. Sci. USA* 79 (1982) 1331–1335.
- [237] J. Newton, Evolutionary game theory: A renaissance, *Games* 9 (2018) 31.

- [238] T. Kay, L. Keller, L. Lehmann, The evolution of altruism and the serial rediscovery of the role of relatedness, *Proc. Natl. Acad. Sci. USA* 117 (2020) 28894–28898.
- [239] S. J. Tudge, M. Brede, A tale of two theorems: Comment on “universal scaling for the dilemma strength in evolutionary games” by Z. Wang et al., *Phys. Life Rev.* 14 (2015) 49–51.
- [240] H. Ohtsuki, C. Hauert, E. Lieberman, M. A. Nowak, A simple rule for the evolution of cooperation on graphs and social networks, *Nature* 441 (2006) 502–505.
- [241] H. Ohtsuki, M. A. Nowak, Evolutionary games on cycles, *Proc. R. Soc. B* 273 (2006) 2249–2256.
- [242] H. Ohtsuki, M. A. Nowak, The replicator equation on graphs, *J. Theor. Biol.* 243 (2006) 86–97.
- [243] M. A. Nowak, R. M. May, Evolutionary games and spatial chaos, *Nature* 359 (1992) 826–829.
- [244] M. A. Nowak, R. M. May, The spatial dilemmas of evolution, *Int. J. Bifurc. Chaos* 3 (1993) 35–78.
- [245] J. Gómez-Gardenes, M. Campillo, L. M. Floría, Y. Moreno, Dynamical organization of cooperation in complex topologies, *Phys. Rev. Lett.* 98 (2007) 108103.
- [246] J. Poncela, J. Gómez-Gardenes, L. M. Floría, Y. Moreno, Robustness of cooperation in the evolutionary prisoner’s dilemma on complex networks, *New J. Physics* 9 (2007) 184.
- [247] J. Gómez-Gardeñes, J. Poncela, L. M. Floría, Y. Moreno, Natural selection of cooperation and degree hierarchy in heterogeneous populations, *J. Theor. Biol.* 253 (2008) 296–301.
- [248] S. Devlin, T. Treloar, Evolution of cooperation through the heterogeneity of random networks, *Phys. Rev. E* 79 (2009) 016107.
- [249] J. Poncela, J. Gómez-Gardeñes, L. M. Floría, Y. Moreno, A. Sánchez, Cooperative scale-free networks despite the presence of defector hubs, *EPL (Europhys. Lett.)* 88 (2009) 38003.

- [250] F. C. Santos, J. M. Pacheco, Scale-free networks provide a unifying framework for the emergence of cooperation, *Phys. Rev. Lett.* 95 (2005) 098104.
- [251] A. Szolnoki, M. Perc, Z. Danku, Towards effective payoffs in the prisoner's dilemma game on scale-free networks, *Physica A* 387 (2008) 2075–2082.
- [252] A. Yamauchi, J. Tanimoto, A. Hagishima, What controls network reciprocity in the Prisoner's Dilemma game?, *BioSystems* 102 (2010) 82–87.
- [253] M. A. Nowak, C. E. Tarnita, T. Antal, Evolutionary dynamics in structured populations, *Philos. Trans. R. Soc. B* 365 (2010) 19–30.
- [254] B. Allen, G. Lippner, Y.-T. Chen, B. Fotouhi, N. Momeni, S.-T. Yau, M. A. Nowak, Evolutionary dynamics on any population structure, *Nature* 544 (2017) 227–230.
- [255] Q. Su, A. Li, L. Zhou, L. Wang, Interactive diversity promotes the evolution of cooperation in structured populations, *New J. Physics* 18 (2016) 103007.
- [256] I. Sendiña-Nadal, I. Leyva, M. Perc, D. Papo, M. Jusup, Z. Wang, J. A. Almendral, P. Manshour, S. Boccaletti, Diverse strategic identities induce dynamical states in evolutionary games, *Phys. Rev. Res.* 2 (2020) 043168.
- [257] Q. Su, A. Li, L. Wang, Evolution of cooperation with interactive identity and diversity, *J. Theor. Biol.* 442 (2018) 149–157.
- [258] D. Jia, X. Wang, Z. Song, I. Romić, X. Li, M. Jusup, Z. Wang, Evolutionary dynamics drives role specialization in a community of players, *J. R. Soc. Interface* 17 (2020) 20200174.
- [259] A. Szolnoki, M. Mobilia, L.-L. Jiang, B. Szczesny, A. M. Rucklidge, M. Perc, Cyclic dominance in evolutionary games: a review, *J. R. Soc. Interface* 11 (2014) 20140735.
- [260] G. Szabó, C. Hauert, Evolutionary prisoner's dilemma games with voluntary participation, *Phys. Rev. E* 66 (2002) 062903.

- [261] C. Shen, M. Jusup, L. Shi, Z. Wang, M. Perc, P. Holme, Exit rights open complex pathways to cooperation, *J. R. Soc. Interface* 18 (2021) 20200777.
- [262] H. Guo, Z. Song, S. Geček, X. Li, M. Jusup, M. Perc, Y. Moreno, S. Boccaletti, Z. Wang, A novel route to cyclic dominance in voluntary social dilemmas, *J. R. Soc. Interface* 17 (2020) 20190789.
- [263] C. Hauert, M. Doebeli, Spatial structure often inhibits the evolution of cooperation in the snowdrift game, *Nature* 428 (2004) 643–646.
- [264] M. Gosak, R. Markovič, J. Dolensek, M. S. Rupnik, M. Marhl, A. Stožer, M. Perc, Network science of biological systems at different scales: A review, *Phys. Life Rev.* 24 (2018) 118–135.
- [265] S. Pilosof, M. A. Porter, M. Pascual, S. Kéfi, The multilayer nature of ecological networks, *Nat. Ecol. Evol.* 1 (2017) 0101.
- [266] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, S. Havlin, Catastrophic cascade of failures in interdependent networks, *Nature* 464 (2010) 1025–1028.
- [267] J. Gao, S. V. Buldyrev, S. Havlin, H. E. Stanley, Robustness of a network of networks, *Phys. Rev. Lett.* 107 (2011) 195701.
- [268] M. J. Pocock, D. M. Evans, J. Memmott, The robustness and restoration of a network of ecological networks, *Science* 335 (2012) 973–977.
- [269] G. Dong, J. Gao, R. Du, L. Tian, H. E. Stanley, S. Havlin, Robustness of network of networks under targeted attack, *Phys. Rev. E* 87 (2013) 052804.
- [270] D. M. Evans, M. J. Pocock, J. Memmott, The robustness of a network of ecological networks to habitat loss, *Ecol. Lett.* 16 (2013) 844–852.
- [271] Z. Wang, L. Wang, A. Szolnoki, M. Perc, Evolutionary games on multilayer networks: a colloquium, *Eur. Phys. J. B* 88 (2015) 1–15.
- [272] M. De Domenico, A. Solé-Ribalta, E. Cozzo, M. Kivela, Y. Moreno, M. A. Porter, S. Gómez, A. Arenas, Mathematical formulation of multilayer networks, *Phys. Rev. X* 3 (2013) 041022.

- [273] M. Kivelä, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, M. A. Porter, Multilayer networks, *J. Complex Netw.* 2 (2014) 203–271.
- [274] J. Gómez-Gardenes, I. Reinares, A. Arenas, L. M. Floría, Evolution of cooperation in multiplex networks, *Sci. Rep.* 2 (2012) 620.
- [275] L.-L. Jiang, M. Perc, Spreading of cooperative behaviour across interdependent groups, *Sci. Rep.* 3 (2013) 2483.
- [276] J. Gómez-Gardenes, C. Gracia-Lázaro, L. M. Floria, Y. Moreno, Evolutionary dynamics on interdependent populations, *Phys. Rev. E* 86 (2012) 056113.
- [277] Z. Wang, A. Szolnoki, M. Perc, Evolution of public cooperation on interdependent networks: The impact of biased utility functions, *EPL (Europhys. Lett.)* 97 (2012) 48001.
- [278] C. Shen, C. Chu, L. Shi, M. Jusup, M. Perc, Z. Wang, Coevolutionary resolution of the public goods dilemma in interdependent structured populations, *EPL (Europhys. Lett.)* 124 (2018) 48003.
- [279] A.-L. Barabasi, The origin of bursts and heavy tails in human dynamics, *Nature* 435 (2005) 207–211.
- [280] A. Vazquez, Exact results for the Barabási model of human dynamics, *Phys. Rev. Lett.* 95 (2005) 248701.
- [281] M. Starnini, A. Baronchelli, R. Pastor-Satorras, Modeling human dynamics of face-to-face interaction networks, *Phys. Rev. Lett.* 110 (2013) 168701.
- [282] P. Holme, J. Saramäki, Temporal networks, *Phys. Rep.* 519 (2012) 97–125.
- [283] P. Holme, Modern temporal network theory: a colloquium, *Eur. Phys. J. B* 88 (2015) 1–30.
- [284] N. Masuda, R. Lambiotte, A guide to temporal networks, World Scientific, 2016.

- [285] N. Masuda, K. Klemm, V. M. Eguíluz, Temporal networks: slowing down diffusion by long lasting interactions, *Phys. Rev. Lett.* 111 (2013) 188701.
- [286] I. Scholtes, N. Wider, R. Pfitzner, A. Garas, C. J. Tessone, F. Schweitzer, Causality-driven slow-down and speed-up of diffusion in non-Markovian temporal networks, *Nat. Commun.* 5 (2014) 1–9.
- [287] N. Masuda, M. A. Porter, R. Lambiotte, Random walks and diffusion on networks, *Phys. Rep.* 716 (2017) 1–58.
- [288] E. Valdano, L. Ferreri, C. Poletto, V. Colizza, Analytical computation of the epidemic threshold on temporal networks, *Phys. Rev. X* 5 (2015) 021005.
- [289] L. E. Rocha, F. Liljeros, P. Holme, Simulated epidemics in an empirical spatiotemporal network of 50,185 sexual contacts, *PLOS Comput. Biol.* 7 (2011) e1001109.
- [290] A. Li, S. P. Cornelius, Y.-Y. Liu, L. Wang, A.-L. Barabási, The fundamental advantages of temporal networks, *Science* 358 (2017) 1042–1046.
- [291] M. Perc, A. Szolnoki, Coevolutionary games—A mini review, *BioSystems* 99 (2010) 109–125.
- [292] M. G. Zimmermann, V. M. Eguíluz, M. San Miguel, Coevolution of dynamical states and interactions in dynamic networks, *Phys. Rev. E* 69 (2004) 065102.
- [293] V. M. Eguíluz, M. G. Zimmermann, C. J. Cela-Conde, M. S. Miguel, Cooperation and the emergence of role differentiation in the dynamics of social networks, *Am. J. Sociol.* 110 (2005) 977–1008.
- [294] F. Fu, T. Wu, L. Wang, Partner switching stabilizes cooperation in coevolutionary prisoner’s dilemma, *Phys. Rev. E* 79 (2009) 036101.
- [295] F. Du, F. Fu, Partner selection shapes the strategic and topological evolution of cooperation, *Dyn. Games Appl.* 1 (2011) 354–369.

- [296] D. Melamed, M. Sweitzer, B. Simpson, J. Z. Abernathy, A. Harrell, C. W. Munn, Homophily and segregation in cooperative networks, *Am. J. Sociol.* 125 (2020) 1084–1127.
- [297] A. Cardillo, G. Petri, V. Nicosia, R. Sinatra, J. Gómez-Gardenes, V. Latora, Evolutionary dynamics of time-resolved social interactions, *Phys. Rev. E* 90 (2014) 052825.
- [298] A. Li, L. Zhou, Q. Su, S. P. Cornelius, Y.-Y. Liu, L. Wang, S. A. Levin, Evolution of cooperation on temporal networks, *Nat. Commun.* 11 (2020) 2259.
- [299] N. Perra, B. Gonçalves, R. Pastor-Satorras, A. Vespignani, Activity driven modeling of time varying networks, *Sci. Rep.* 2 (2012) 469.
- [300] U. Alvarez-Rodriguez, F. Battiston, G. F. de Arruda, Y. Moreno, M. Perc, V. Latora, Evolutionary dynamics of higher-order interactions in social networks, *Nat. Hum. Behav.* 5 (2021) 586–595.
- [301] F. C. Santos, M. D. Santos, J. M. Pacheco, Social diversity promotes the emergence of cooperation in public goods games, *Nature* 454 (2008) 213–216.
- [302] A. Szolnoki, M. Perc, G. Szabó, Topology-independent impact of noise on cooperation in spatial public goods games, *Phys. Rev. E* 80 (2009) 056109.
- [303] F. Battiston, G. Cencetti, I. Iacopini, V. Latora, M. Lucas, A. Pataña, J.-G. Young, G. Petri, Networks beyond pairwise interactions: structure and dynamics, *Phys. Rep.* 874 (2020) 1–92.
- [304] C. Hauert, G. Szabo, Prisoner’s dilemma and public goods games in different geometries: compulsory versus voluntary interactions, *Complexity* 8 (2003) 31–38.
- [305] L. Wu, D. Wang, J. A. Evans, Large teams develop and small teams disrupt science and technology, *Nature* 566 (2019) 378–382.
- [306] C. P. Roca, J. A. Cuesta, A. Sánchez, Evolutionary game theory: Temporal and spatial effects beyond replicator dynamics, *Phys. Life Rev.* 6 (2009) 208–249.

- [307] M. A. Nowak, A. Sasaki, C. Taylor, D. Fudenberg, Emergence of cooperation and evolutionary stability in finite populations, *Nature* 428 (2004) 646–650.
- [308] L. Hindersin, B. Wu, A. Traulsen, J. García, Computation and simulation of evolutionary game dynamics in finite populations, *Sci. Rep.* 9 (2019) 6946.
- [309] E. Lieberman, C. Hauert, M. A. Nowak, Evolutionary dynamics on graphs, *Nature* 433 (2005) 312–316.
- [310] F. C. Santos, J. M. Pacheco, T. Lenaerts, Evolutionary dynamics of social dilemmas in structured heterogeneous populations, *Proc. Natl. Acad. Sci. USA* 103 (2006) 3490–3494.
- [311] W. Maciejewski, F. Fu, C. Hauert, Evolutionary game dynamics in populations with heterogeneous structures, *PLOS Comput. Biol.* 10 (2014) e1003567.
- [312] B. Wu, C. S. Gokhale, M. van Veelen, L. Wang, A. Traulsen, Interpretations arising from Wrightian and Malthusian fitness under strong frequency dependent selection, *Ecol. Evol.* 3 (2013) 1276–1280.
- [313] M. Doebeli, Y. Ispolatov, B. Simon, Point of view: Towards a mechanistic foundation of evolutionary theory, *eLife* 6 (2017) e23804.
- [314] B. Wu, J. García, C. Hauert, A. Traulsen, Extrapolating weak selection in evolutionary games, *PLOS Comput. Biol.* 9 (2013) e1003381.
- [315] B. Wu, B. Bauer, T. Galla, A. Traulsen, Fitness-based models and pairwise comparison models of evolutionary games are typically different—even in unstructured populations, *New J. Phys.* 17 (2015) 023043.
- [316] A. Sánchez, Theory must be informed by experiments (and back): comment on “Universal scaling for the dilemma strength in evolutionary games”, by Z. Wang et al., *Phys. Life Rev.* 14 (2015) 52–53.
- [317] A. Sánchez, Physics of human cooperation: experimental evidence and theoretical models, *J. Stat. Mech. Theory Exp.* 2018 (2018) 024001.

- [318] A. Dreber, D. G. Rand, D. Fudenberg, M. A. Nowak, Winners don't punish, *Nature* 452 (2008) 348–351.
- [319] J.-J. Wu, B.-Y. Zhang, Z.-X. Zhou, Q.-Q. He, X.-D. Zheng, R. Cressman, Y. Tao, Costly punishment does not always increase cooperation, *Proc. Natl. Acad. Sci. USA* 106 (2009) 17448–17451.
- [320] X. Li, M. Jusup, Z. Wang, H. Li, L. Shi, B. Podobnik, H. E. Stanley, S. Havlin, S. Boccaletti, Punishment diminishes the benefits of network reciprocity in social dilemma experiments, *Proc. Natl. Acad. Sci. USA* 115 (2018) 30–35.
- [321] C. Gracia-Lázaro, A. Ferrer, G. Ruiz, A. Tarancón, J. A. Cuesta, A. Sánchez, Y. Moreno, Heterogeneous networks do not promote cooperation when humans play a Prisoner's Dilemma, *Proc. Natl. Acad. Sci. USA* 109 (2012) 12922–12926.
- [322] J. Grujić, C. Gracia-Lázaro, M. Milinski, D. Semmann, A. Traulsen, J. A. Cuesta, Y. Moreno, A. Sánchez, A comparative analysis of spatial Prisoner's Dilemma experiments: Conditional cooperation and payoff irrelevance, *Sci. Rep.* 4 (2014) 4615.
- [323] D. G. Rand, M. A. Nowak, J. H. Fowler, N. A. Christakis, Static network structure can stabilize human cooperation, *Proc. Natl. Acad. Sci. USA* 111 (2014) 17093–17098.
- [324] G. Szabó, C. Tóke, Evolutionary prisoner's dilemma game on a square lattice, *Phys. Rev. E* 58 (1998) 69.
- [325] A. Traulsen, M. A. Nowak, J. M. Pacheco, Stochastic dynamics of invasion and fixation, *Phys. Rev. E* 74 (2006) 011909.
- [326] J. Apesteguia, S. Huck, J. Oechssler, Imitation—theory and experimental evidence, *J. Econ. Theory* 136 (2007) 217–235.
- [327] A. Traulsen, D. Semmann, R. D. Sommerfeld, H.-J. Krambeck, M. Milinski, Human strategy updating in evolutionary games, *Proc. Natl. Acad. Sci. USA* 107 (2010) 2962–2966.
- [328] J. Grujić, T. Lenaerts, Do people imitate when making decisions? Evidence from a spatial Prisoner's Dilemma experiment, *R. Soc. Open Sci.* 7 (2020) 200618.

- [329] J. Poncela-Casasnovas, M. Gutiérrez-Roig, C. Gracia-Lázaro, J. Vicens, J. Gómez-Gardeñes, J. Perelló, Y. Moreno, J. Duch, A. Sánchez, Humans display a reduced set of consistent behavioral phenotypes in dyadic games, *Sci. Adv.* 2 (2016) e1600451.
- [330] A. Peysakhovich, M. A. Nowak, D. G. Rand, Humans display a ‘cooperative phenotype’ that is domain general and temporally stable, *Nat. Commun.* 5 (2014) 4939.
- [331] R. Axelrod, W. D. Hamilton, The evolution of cooperation, *Science* 211 (1981) 1390–1396.
- [332] J. Bendor, P. Swistak, The evolutionary stability of cooperation, *Am. Political Sci. Rev.* 91 (1997) 290–307.
- [333] C. Keser, F. Van Winden, Conditional cooperation and voluntary contributions to public goods, *Scand. J. Econ.* 102 (2000) 23–39.
- [334] U. Fischbacher, S. Gächter, E. Fehr, Are people conditionally cooperative? Evidence from a public goods experiment, *Econ. Lett.* 71 (2001) 397–404.
- [335] J. Grujić, C. Fosco, L. Araujo, J. A. Cuesta, A. Sánchez, Social experiments in the mesoscale: Humans playing a spatial prisoner’s dilemma, *PLOS ONE* 5 (2010) e13749.
- [336] L. Shi, I. Romić, Y. Ma, Z. Wang, B. Podobnik, H. E. Stanley, P. Holme, M. Jusup, Freedom of choice adds value to public goods, *Proc. Natl. Acad. Sci. USA* 117 (2020) 17516–17521.
- [337] M. Gutiérrez-Roig, C. Gracia-Lázaro, J. Perelló, Y. Moreno, A. Sánchez, Transition from reciprocal cooperation to persistent behaviour in social dilemmas at the end of adolescence, *Nat. Commun.* 5 (2014) 4362.
- [338] Y. Horita, M. Takezawa, K. Inukai, T. Kita, N. Masuda, Reinforcement learning accounts for moody conditional cooperation behavior: experimental results, *Sci. Rep.* 7 (2017) 1–10.
- [339] A. Tversky, Elimination by aspects: A theory of choice., *Psychol. Rev.* 79 (1972) 281–299.

- [340] D. Kahneman, A. Tversky, Prospect theory: An analysis of decision under risk, *Econometrica* 47 (1979) 363–391.
- [341] P. J. Schoemaker, The expected utility model: Its variants, purposes, evidence and limitations, *J. Econ. Lit.* 20 (1982) 529–563.
- [342] Z. Wang, M. Jusup, R.-W. Wang, L. Shi, Y. Iwasa, Y. Moreno, J. Kurths, Onymity promotes cooperation in social dilemma experiments, *Sci. Adv.* 3 (2017) e1601444.
- [343] E. Hoffman, K. McCabe, V. L. Smith, Social distance and other-regarding behavior in dictator games, *Am. Econ. Rev.* 86 (1996) 653–660.
- [344] Z. Wang, M. Jusup, L. Shi, J.-H. Lee, Y. Iwasa, S. Boccaletti, Exploiting a cognitive bias promotes cooperation in social dilemma experiments, *Nat. Commun.* 9 (2018) 2954.
- [345] J. C. Pettibone, D. H. Wedell, Examining models of nondominated decoy effects across judgment and choice, *Organ. Behav. Hum. Decis. Process.* 81 (2000) 300–328.
- [346] T. M. A. D. Group, T. W. Fawcett, B. Fallenstein, A. D. Higginson, A. I. Houston, D. E. Mallpress, P. C. Trimmer, J. M. McNamara, The evolution of decision rules in complex environments, *Trends Cogn. Sci.* 18 (2014) 153–161.
- [347] D. G. Rand, S. Arbesman, N. A. Christakis, Dynamic social networks promote cooperation in experiments with humans, *Proc. Natl. Acad. Sci. USA* 108 (2011) 19193–19198.
- [348] A. Gardner, S. A. West, Cooperation and punishment, especially in humans, *Am. Nat.* 164 (2004) 753–764.
- [349] C. Hauert, A. Traulsen, H. Brandt, M. A. Nowak, K. Sigmund, Via freedom to coercion: the emergence of costly punishment, *Science* 316 (2007) 1905–1907.
- [350] A. P. Melis, D. Semmann, How is human cooperation different?, *Philos. Trans. R. Soc. B* 365 (2010) 2663–2674.

- [351] N. Nikiforakis, Punishment and counter-punishment in public good games: Can we really govern ourselves?, *J. Public Econ.* 92 (2008) 91–112.
- [352] N. Nikiforakis, D. Engelmann, Altruistic punishment and the threat of feuds, *J. Econ. Behav. Organ.* 78 (2011) 319–332.
- [353] L. Balafoutas, K. Grechenig, N. Nikiforakis, Third-party punishment and counter-punishment in one-shot interactions, *Econ. Lett.* 122 (2014) 308–310.
- [354] K. Fehl, D. J. van der Post, D. Semmann, Co-evolution of behaviour and social network structure promotes human cooperation, *Ecol. Lett.* 14 (2011) 546–551.
- [355] J. W. Pepper, B. B. Smuts, A mechanism for the evolution of altruism among nonkin: positive assortment through environmental feedback, *Am. Nat.* 160 (2002) 205–213.
- [356] J. A. Fletcher, M. Zwick, Unifying the theories of inclusive fitness and reciprocal altruism, *Am. Nat.* 168 (2006) 252–262.
- [357] J. A. Fletcher, M. Doebeli, A simple and general explanation for the evolution of altruism, *Proc. R. Soc. B* 276 (2009) 13–19.
- [358] B. Momeni, A. J. Waite, W. Shou, Spatial self-organization favors heterotypic cooperation over cheating, *eLife* 2 (2013) e00960.
- [359] R. A. Klein, K. A. Ratliff, M. Vianello, R. B. Adams Jr, Š. Bahník, M. J. Bernstein, K. Bocian, M. J. Brandt, B. Brooks, C. C. Brumbaugh, et al., Investigating variation in replicability, *Soc. Psychol.* 45 (2014) 142–152.
- [360] O. S. Collaboration, Estimating the reproducibility of psychological science, *Science* 349 (2015) aac4716.
- [361] C. F. Camerer, A. Dreber, F. Holzmeister, T.-H. Ho, J. Huber, M. Johannesson, M. Kirchler, G. Nave, B. A. Nosek, T. Pfeiffer, et al., Evaluating the replicability of social science experiments in *Nature* and *Science* between 2010 and 2015, *Nat. Hum. Behav.* 2 (2018) 637–644.

- [362] R. A. Klein, M. Vianello, F. Hasselman, B. G. Adams, R. B. Adams Jr, S. Alper, M. Aveyard, J. R. Axt, M. T. Babalola, Š. Bahník, et al., Many Labs 2: Investigating variation in replicability across samples and settings, *Adv. Methods Pract. Psychol. Sci.* 1 (2018) 443–490.
- [363] M. Serra-Garcia, U. Gneezy, Nonreplicable publications are cited more than replicable ones, *Sci. Adv.* 7 (2021) eabd1705.
- [364] B. A. Nosek, G. Alter, G. C. Banks, D. Borsboom, S. D. Bowman, S. J. Breckler, S. Buck, C. D. Chambers, G. Chin, G. Christensen, et al., Promoting an open research culture, *Science* 348 (2015) 1422–1425.
- [365] M. R. Munafò, B. A. Nosek, D. V. M. Bishop, K. S. Button, C. D. Chambers, N. P. Du Sert, U. Simonsohn, E.-J. Wagenmakers, J. J. Ware, J. P. A. Ioannidis, A manifesto for reproducible science, *Nat. Hum. Behav* 1 (2017) 0021.
- [366] P. E. Shrout, J. L. Rodgers, Psychology, science, and knowledge construction: Broadening perspectives from the replication crisis, *Annu. Rev. Psychol.* 69 (2018) 487–510.
- [367] D. J. Benjamin, J. O. Berger, M. Johannesson, B. A. Nosek, E.-J. Wagenmakers, R. Berk, K. A. Bollen, B. Brembs, L. Brown, C. Camerer, et al., Redefine statistical significance, *Nat. Hum. Behav.* 2 (2018) 6–10.
- [368] M. J. Brandt, H. IJzerman, A. Dijksterhuis, F. J. Farach, J. Geller, R. Giner-Sorolla, J. A. Grange, M. Perugini, J. R. Spies, A. Van’t Veer, The replication recipe: What makes for a convincing replication?, *J. Exp. Soc. Psychol.* 50 (2014) 217–224.
- [369] S. E. Maxwell, M. Y. Lau, G. S. Howard, Is psychology suffering from a replication crisis? What does “failure to replicate” really mean?, *Am Psychol.* 70 (2015) 487–498.
- [370] B. A. Nosek, C. R. Ebersole, A. C. DeHaven, D. T. Mellor, The preregistration revolution, *Proc. Natl. Acad. Sci. USA* 115 (2018) 2600–2606.

- [371] B. A. Nosek, E. D. Beck, L. Campbell, J. K. Flake, T. E. Hardwicke, D. T. Mellor, A. E. van't Veer, S. Vazire, Preregistration is hard, and worthwhile, *Trends Cogn. Sci.* 23 (2019) 815–818.
- [372] C. K. Soderberg, T. M. Errington, S. R. Schiavone, J. Bottesini, F. S. Thorn, S. Vazire, K. M. Esterling, B. A. Nosek, Initial evidence of research quality of registered reports compared with the standard publishing model, *Nat. Hum. Behav.* 5 (2021) 990–997.
- [373] M. Muthukrishna, J. Henrich, A problem in theory, *Nat. Hum. Behav.* 3 (2019) 221–229.
- [374] G. Brumfiel, Particles break light-speed limit, 2011. Available at: <https://www.nature.com/articles/news.2011.554>.
- [375] T. Adam, N. Agafonova, A. Aleksandrov, O. Altinok, P. A. Sanchez, A. Anokhina, S. Aoki, A. Ariga, T. Ariga, D. Autiero, et al., Measurement of the neutrino velocity with the OPERA detector in the CNGS beam, *J. High Energy Phys.* 2012 (2012) 1–37.
- [376] J.-H. Lee, K. Sigmund, U. Dieckmann, Y. Iwasa, Games of corruption: How to suppress illegal logging, *J. Theor. Biol.* 367 (2015) 1–13.
- [377] J.-H. Lee, Y. Iwasa, U. Dieckmann, K. Sigmund, Social evolution leads to persistent corruption, *Proc. Natl. Acad. Sci.* 116 (2019) 13276–13281.
- [378] M. R. Arefin, T. Masaki, K. A. Kabir, J. Tanimoto, Interplay between cost and effectiveness in influenza vaccine uptake: a vaccination game approach, *Proc. R. Soc. A* 475 (2019) 20190608.
- [379] K. A. Kabir, M. Jusup, J. Tanimoto, Behavioral incentives in a vaccination-dilemma setting with optional treatment, *Phys. Rev. E* 100 (2019) 062402.
- [380] M. Nowak, K. Sigmund, A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game, *Nature* 364 (1993) 56–58.
- [381] H. Ohtsuki, Y. Iwasa, The leading eight: social norms that can maintain cooperation by indirect reciprocity, *J. Theor. Biol.* 239 (2006) 435–444.

- [382] H. Ohtsuki, Y. Iwasa, Global analyses of evolutionary dynamics and exhaustive search for social norms that maintain cooperation by reputation, *J. Theor. Biol.* 244 (2007) 518–531.
- [383] P. Holme, Network dynamics of ongoing social relationships, *EPL (Europhys. Lett.)* 64 (2003) 427.
- [384] M. E. Newman, J. Park, Why social networks are different from other types of networks, *Phys. Rev. E* 68 (2003) 036122.
- [385] P. Holme, C. R. Edling, F. Liljeros, Structure and time evolution of an Internet dating community, *Soc. Netw.* 26 (2004) 155–174.
- [386] L. E. Rocha, F. Liljeros, P. Holme, Information dynamics shape the sexual networks of Internet-mediated prostitution, *Proc. Natl. Acad. Sci. USA* 107 (2010) 5706–5711.
- [387] N. Pinter-Wollman, E. A. Hobson, J. E. Smith, A. J. Edelman, D. Shizuka, S. De Silva, J. S. Waters, S. D. Prager, T. Sasaki, G. Wittemyer, et al., The dynamics of animal social networks: analytical, conceptual, and theoretical advances, *Behav. Ecol.* 25 (2014) 242–255.
- [388] M. Jalili, Y. Orouskhani, M. Asgari, N. Alipourfard, M. Perc, Link prediction in multiplex online social networks, *R. Soc. Open Sci.* 4 (2017) 160863.
- [389] S. Battiston, M. Puliga, R. Kaushik, P. Tasca, G. Caldarelli, Debtrank: Too central to fail? Financial networks, the FED and systemic risk, *Sci. Rep.* 2 (2012) 541.
- [390] L. E. C. Rocha, P. Holme, The network organisation of consumer complaints, *EPL (Europhys. Lett.)* 91 (2010) 28005.
- [391] T. Kito, A. Brintrup, S. New, F. Reed-Tsochas, The structure of the Toyota supply network: an empirical analysis, 2014. Social Science Research Network (SSRN) Paper No. 2412512. Available at: <https://ssrn.com/abstract=2412512>.
- [392] H. Inoue, Y. Todo, Firm-level propagation of shocks through supply-chain networks, *Nat. Sustain.* 2 (2019) 841–847.

- [393] P. Holme, Congestion and centrality in traffic flow on complex networks, *Adv. Complex Syst.* 6 (2003) 163–176.
- [394] D. DeLaurentis, E.-P. Han, T. Kotegawa, Network-theoretic approach for analyzing connectivity in air transportation networks, *J. Aircr.* 45 (2008) 1669–1679.
- [395] Y. Hu, D. Zhu, Empirical analysis of the worldwide maritime transportation network, *Physica A* 388 (2009) 2061–2071.
- [396] F. Xie, D. Levinson, Modeling the growth of transportation networks: A comprehensive review, *Netw. Spat. Econ.* 9 (2009) 291–307.
- [397] S. Derrible, C. Kennedy, Applications of graph theory and network science to transit network design, *Transp. Rev.* 31 (2011) 495–519.
- [398] W.-B. Du, X.-L. Zhou, M. Jusup, Z. Wang, Physics of transportation: Towards optimal capacity using the multilayer network framework, *Sci. Rep.* 6 (2016) 19059.
- [399] M. Rosas-Casals, S. Valverde, R. V. Solé, Topological vulnerability of the European power grid under errors and attacks, *Int. J. Bifurc. Chaos* 17 (2007) 2465–2475.
- [400] G. A. Paganì, M. Aiello, The power grid as a complex network: a survey, *Physica A* 392 (2013) 2688–2700.
- [401] H. Kim, S. H. Lee, P. Holme, Community consistency determines the stability transition window of power-grid nodes, *New J. Phys.* 17 (2015) 113005.
- [402] J. F. Donges, Y. Zou, N. Marwan, J. Kurths, Complex networks in climate dynamics, *Eur. Phys. J. Spec. Top.* 174 (2009) 157–179.
- [403] J. Fan, J. Meng, J. Ludescher, X. Chen, Y. Ashkenazy, J. Kurths, S. Havlin, H. J. Schellnhuber, Statistical physics approaches to the complex earth system, *Phys. Rep.* 869 (2021) 1–84.
- [404] A.-L. Barabási, Network medicine—from obesity to the “diseasome”, *N. Engl. J. Med.* 357 (2007) 404–407.

- [405] S. G. Hofmann, J. Curtiss, A complex network approach to clinical science, *Eur. J. Clin. Invest.* 48 (2018) e12986.
- [406] A.-L. Barabási, G. Menichetti, J. Loscalzo, The unmapped chemical complexity of our diet, *Nat. Food* 1 (2020) 33–37.
- [407] J. C. S. Herrera, The contribution of network science to the study of food recipes. A review paper, *Appetite* 159 (2020) 105048.
- [408] J. M. Buldú, J. Busquets, J. H. Martínez, J. L. Herrera-Diestra, I. Echegoyen, J. Galeano, J. Luque, Using network science to analyse football passing networks: Dynamics, space, time, and the multilayer nature of the game, *Front. Psychol.* 9 (2018) 1900.
- [409] J. Buldú, J. Busquets, I. Echegoyen, et al., Defining a historic football team: Using Network Science to analyze Guardiola’s FC Barcelona, *Sci. Rep.* 9 (2019) 13602.
- [410] J. D. Noh, H. Rieger, Random walks on complex networks, *Phys. Rev. Lett.* 92 (2004) 118701.
- [411] T. J. Perkins, E. Foxall, L. Glass, R. Edwards, A scaling law for random walks on networks, *Nat. Commun.* 5 (2014) 5121.
- [412] C. W. Wu, Synchronization in networks of nonlinear dynamical systems coupled via a directed graph, *Nonlinearity* 18 (2005) 1057.
- [413] L. Scardovi, R. Sepulchre, Synchronization in networks of identical linear systems, *Automatica* 45 (2009) 2557–2562.
- [414] C. I. Del Genio, J. Gómez-Gardeñes, I. Bonamassa, S. Boccaletti, Synchronization in networks with multiple interaction layers, *Sci. Adv.* 2 (2016) e1601679.
- [415] S. Bansal, B. T. Grenfell, L. A. Meyers, When individual behaviour matters: homogeneous and network models in epidemiology, *J. R. Soc. Interface* 4 (2007) 879–891.
- [416] L. Danon, A. P. Ford, T. House, C. P. Jewell, M. J. Keeling, G. O. Roberts, J. V. Ross, M. C. Vernon, Networks and the epidemiology of infectious disease, *Interdiscip. Perspect. Infect. Dis.* 2011 (2011) 284909.

- [417] D. Brockmann, D. Helbing, The hidden geometry of complex, network-driven contagion phenomena, *science* 342 (2013) 1337–1342.
- [418] N. P. Hummon, P. Doreian, Some dynamics of social balance processes: bringing Heider back into balance theory, *Soc. Netw.* 25 (2003) 17–49.
- [419] T. Antal, P. L. Krapivsky, S. Redner, Dynamics of social balance on networks, *Phys. Rev. E* 72 (2005) 036121.
- [420] T. Antal, P. L. Krapivsky, S. Redner, Social balance on networks: The dynamics of friendship and enmity, *Physica D* 224 (2006) 130–136.
- [421] I. Iacopini, S. Milojević, V. Latora, Network dynamics of innovation processes, *Phys. Rev. Lett.* 120 (2018) 048301.
- [422] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, D.-U. Hwang, Complex networks: Structure and dynamics, *Phys. Rep.* 424 (2006) 175–308.
- [423] S. Fortunato, Community detection in graphs, *Phys Rep.* 486 (2010) 75–174.
- [424] S. Fortunato, D. Hric, Community detection in networks: A user guide, *Phys. Rep.* 659 (2016) 1–44.
- [425] M. Barthélemy, Spatial networks, *Phys. Rep.* 499 (2011) 1–101.
- [426] S. Boccaletti, G. Bianconi, R. Criado, C. I. Del Genio, J. Gómez-Gardenes, M. Romance, I. Sendina-Nadal, Z. Wang, M. Zanin, The structure and dynamics of multilayer networks, *Phys. Rep.* 544 (2014) 1–122.
- [427] G. F. de Arruda, F. A. Rodrigues, Y. Moreno, Fundamentals of spreading processes in single and multilayer complex networks, *Phys. Rep.* 756 (2018) 1–59.
- [428] A. Lancichinetti, S. Fortunato, Community detection algorithms: a comparative analysis, *Phys. Rev. E* 80 (2009) 056117.

- [429] A. Lancichinetti, F. Radicchi, J. J. Ramasco, S. Fortunato, Finding statistically significant communities in networks, *PLOS ONE* 6 (2011) e18961.
- [430] D. Hric, R. K. Darst, S. Fortunato, Community detection in networks: Structural communities versus ground truth, *Phys. Rev. E* 90 (2014) 062805.
- [431] M. T. Schaub, J.-C. Delvenne, M. Rosvall, R. Lambiotte, The many facets of community detection in complex networks, *Appl. Netw. Sci.* 2 (2017) 4.
- [432] M. Rosvall, J.-C. Delvenne, M. T. Schaub, R. Lambiotte, Different approaches to community detection, in: P. Doreian, V. Batagelj, A. Ferligoj (Eds.), *Advances in network clustering and blockmodeling*, Wiley Online Library, 2019, pp. 105–119.
- [433] K. Chung, Y. Baek, D. Kim, M. Ha, H. Jeong, Generalized epidemic process on modular networks, *Phys. Rev. E* 89 (2014) 052811.
- [434] L. D. Valdez, L. A. Braunstein, S. Havlin, Epidemic spreading on modular networks: The fear to declare a pandemic, *Phys. Rev. E* 101 (2020) 032309.
- [435] N. Masuda, Immunization of networks with community structure, *New J. Phys.* 11 (2009) 123018.
- [436] B. Gross, S. Havlin, Epidemic spreading and control strategies in spatial modular network, *Appl. Netw. Sci.* 5 (2020) 95.
- [437] P. Erdos, On random graphs i, *Publ. Math. Debr.* 6 (1959) 290–297.
- [438] D. J. Watts, S. H. Strogatz, Collective dynamics of ‘small-world’ networks, *Nature* 393 (1998) 440–442.
- [439] R. Albert, A.-L. Barabási, Statistical mechanics of complex networks, *Rev. Mod. Phys.* 74 (2002) 47.
- [440] P. W. Holland, K. B. Laskey, S. Leinhardt, Stochastic blockmodels: First steps, *Soc. Netw.* 5 (1983) 109–137.

- [441] K. Nowicki, T. A. B. Snijders, Estimation and prediction for stochastic blockstructures, *J. Am. Stat. Assoc.* 96 (2001) 1077–1087.
- [442] E. M. Airoldi, D. M. Blei, S. E. Fienberg, E. P. Xing, Mixed membership stochastic blockmodels, *J. Mach. Learn. Res.* 9 (2008) 1981–2014.
- [443] X. Fan, R. Y. D. Xu, L. Cao, Copula mixed-membership stochastic block model, in: G. Brewka (Ed.), *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, Association for Computing Machinery*, 2016, pp. 1462–1468.
- [444] S. Pal, M. Coates, Scalable MCMC in degree corrected stochastic block model, in: S. Sanei, L. Hanzo (Eds.), *2019 IEEE International Conference on Acoustics, Speech and Signal Processing, IEEE*, 2019, pp. 5461–5465.
- [445] C. Aicher, A. Z. Jacobs, A. Clauset, Learning latent block structure in weighted networks, *J. Complex Netw.* 3 (2015) 221–248.
- [446] T. P. Peixoto, Nonparametric weighted stochastic block models, *Phys. Rev. E* 97 (2018) 012306.
- [447] T. P. Peixoto, Inferring the mesoscale structure of layered, edge-valued, and time-varying networks, *Phys. Rev. E* 92 (2015) 042807.
- [448] N. Stanley, S. Shai, D. Taylor, P. J. Mucha, Clustering network layers with the strata multilayer stochastic block model, *IEEE Trans. Netw. Sci. Eng.* 3 (2016) 95–105.
- [449] S. Paul, Y. Chen, Consistent community detection in multi-relational data through restricted multi-layer stochastic blockmodel, *Electron. J. Stat.* 10 (2016) 3807–3870.
- [450] T. Valles-Catala, F. A. Massucci, R. Guimera, M. Sales-Pardo, Multilayer stochastic block models reveal the multilayer structure of complex networks, *Phys. Rev. X* 6 (2016) 011036.
- [451] C. De Bacco, E. A. Power, D. B. Larremore, C. Moore, Community detection, link prediction, and layer interdependence in multilayer networks, *Phys. Rev. E* 95 (2017) 042317.

- [452] W. Fu, L. Song, E. P. Xing, Dynamic mixed membership blockmodel for evolving networks, in: A. Danyluk, L. Bottou, M. Littman (Eds.), Proceedings of the 26th Annual International Conference on Machine Learning, Association for Computing Machinery, 2009, pp. 329–336.
- [453] K. S. Xu, A. O. Hero, Dynamic stochastic blockmodels for time-evolving social networks, *IEEE J. Sel. Top. Signal Process.* 8 (2014) 552–562.
- [454] L. Peel, A. Clauset, Detecting change points in the large-scale structure of evolving networks, in: B. Bonet, S. Koenig (Eds.), Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, Association for the Advancement of Artificial Intelligence, 2015, pp. 2914–2920.
- [455] M. Corneli, P. Latouche, F. Rossi, Exact ICL maximization in a non-stationary temporal extension of the stochastic block model for dynamic networks, *Neurocomputing* 192 (2016) 81–91.
- [456] A. Ghasemian, P. Zhang, A. Clauset, C. Moore, L. Peel, Detectability thresholds and optimal algorithms for community structure in dynamic networks, *Phys. Rev. X* 6 (2016) 031005.
- [457] C. Matias, V. Miele, Statistical clustering of temporal networks through a dynamic stochastic block model, *J. R. Stat. Soc. Series B Stat. Methodol.* 79 (2017) 1119–1141.
- [458] T. P. Peixoto, M. Rosvall, Modelling sequences and temporal networks with dynamic community structures, *Nat. Commun.* 8 (2017) 582.
- [459] X. Zhang, C. Moore, M. E. J. Newman, Random graph models for dynamic networks, *Eur. Phys. J. B* 90 (2017) 1–14.
- [460] L. Peel, Active discovery of network roles for predicting the classes of network nodes, *J. Complex Netw.* 3 (2015) 431–449.
- [461] D. Hric, T. P. Peixoto, S. Fortunato, Network structure, metadata, and the prediction of missing nodes and annotations, *Phys. Rev. X* 6 (2016) 031038.

- [462] M. E. J. Newman, A. Clauset, Structure and inference in annotated networks, *Nat. Commun.* 7 (2016) 11863.
- [463] T. P. Peixoto, Bayesian stochastic blockmodeling, in: P. Doreian, V. Batagelj, A. Ferligoj (Eds.), *Advances in network clustering and blockmodeling*, Wiley Online Library, 2019, pp. 289–332.
- [464] T. Funke, T. Becker, Stochastic block models: A comparison of variants and inference methods, *PLOS ONE* 14 (2019) e0215296.
- [465] H. C. White, S. A. Boorman, R. L. Breiger, Social structure from multiple networks. I. Blockmodels of roles and positions, *Am. J. Sociol.* 81 (1976) 730–780.
- [466] P. Arabie, S. A. Boorman, P. R. Levitt, Constructing blockmodels: How and why, *J. Math. Psychol.* 17 (1978) 21–63.
- [467] S. Wasserman, C. Anderson, Stochastic a posteriori blockmodels: Construction and assessment, *Soc. Netw.* 9 (1987) 1–36.
- [468] T. P. Peixoto, Nonparametric Bayesian inference of the microcanonical stochastic block model, *Phys. Rev. E* 95 (2017) 012317.
- [469] B. Karrer, M. E. Newman, Stochastic blockmodels and community structure in networks, *Phys. Rev. E* 83 (2011) 016107.
- [470] L. Gulikers, M. Lelarge, L. Massoulié, A spectral method for community detection in moderately sparse degree-corrected stochastic block models, *Adv. Appl. Probab.* 49 (2017) 686–721.
- [471] Y. Chen, X. Li, J. Xu, Convexified modularity maximization for degree-corrected stochastic block models, *Ann. Stat.* 46 (2018) 1573–1602.
- [472] M. E. Newman, M. Girvan, Finding and evaluating community structure in networks, *Phys. Rev. E* 69 (2004) 026113.
- [473] M. E. Newman, Equivalence between modularity optimization and maximum likelihood methods for community detection, *Phys. Rev. E* 94 (2016) 052315.

- [474] L. Zhang, T. P. Peixoto, Statistical inference of assortative community structures, *Phys. Rev. Res.* 2 (2020) 043271.
- [475] T. P. Peixoto, Entropy of stochastic blockmodel ensembles, *Phys. Rev. E* 85 (2012) 056122.
- [476] J. Rissanen, Modeling by shortest data description, *Automatica* 14 (1978) 465–471.
- [477] M. Rosvall, C. T. Bergstrom, An information-theoretic framework for resolving community structure in complex networks, *Proc. Natl. Acad. Sci.* 104 (2007) 7327–7331.
- [478] T. P. Peixoto, Parsimonious module inference in large networks, *Phys. Rev. Lett.* 110 (2013) 148701.
- [479] T. P. Peixoto, Efficient Monte Carlo and greedy heuristic for the inference of stochastic block models, *Phys. Rev. E* 89 (2014) 012804.
- [480] T. P. Peixoto, Hierarchical block structures and high-resolution model selection in large networks, *Phys. Rev. X* 4 (2014) 011047.
- [481] T. P. Peixoto, Model selection and hypothesis testing for large-scale network models with overlapping groups, *Phys. Rev. X* 5 (2015) 011033.
- [482] C. Lee, D. J. Wilkinson, A review of stochastic block models and extensions for graph clustering, *Appl. Netw. Sci.* 4 (2019) 122.
- [483] D. Jin, Z. Yu, P. Jiao, S. Pan, D. He, J. Wu, P. S. Yu, W. Zhang, A survey of community detection approaches: From statistical modeling to deep learning, 2021. e-print arXiv:2101.01669.
- [484] E. Abbe, Community detection and stochastic block models: recent developments, *J. Mach. Learn. Res.* 18 (2017) 6446–6531.
- [485] M. E. J. Newman, *Networks: An introduction*, Oxford University Press, 2010.
- [486] B. K. Fosdick, D. B. Larremore, J. Nishimura, J. Ugander, Configuring random graph models with fixed degree sequences, *SIAM Review* 60 (2018) 315–355.

- [487] G. Bianconi, Entropy of network ensembles, *Phys. Rev. E* 79 (2009) 036114.
- [488] T. Squartini, J. de Mol, F. den Hollander, D. Garlaschelli, Breaking of ensemble equivalence in networks, *Phys. Rev. Lett.* 115 (2015) 268701.
- [489] E. P. Xing, W. Fu, L. Song, A state-space mixed membership block-model for dynamic network tomography, *Ann. Appl. Stat.* 4 (2010) 535–566.
- [490] X. Yan, Bayesian model selection of stochastic block models, in: V. S. S. Subrahmanian, J. Rokne, R. Kumar, J. Caverlee, H. Tong (Eds.), *Proceedings of the 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, IEEE, 2016, pp. 323–328.
- [491] Y. X. R. Wang, P. J. Bickel, Likelihood-based model selection for stochastic block models, *Ann. Stat.* 45 (2017) 500–528.
- [492] J. Hu, H. Qin, T. Yan, Y. Zhao, Corrected Bayesian information criterion for stochastic block models, *J. Am. Stat. Assoc.* 115 (2020) 1771–1783.
- [493] M. E. J. Newman, G. Reinert, Estimating the number of communities in a network, *Phys. Rev. Lett.* 117 (2016) 078301.
- [494] E. Côme, P. Latouche, Model selection and clustering in stochastic block models based on the exact integrated complete data likelihood, *Stat. Model.* 15 (2015) 564–589.
- [495] W. K. Hastings, Monte carlo sampling methods using Markov chains and their applications, *Biometrika* 57 (1970) 97–109.
- [496] S. Chib, E. Greenberg, Understanding the Metropolis-Hastings algorithm, *Am. Stat.* 49 (1995) 327–335.
- [497] P. K. Gopalan, S. Gerrish, M. Freedman, D. Blei, D. Mimno, Scalable inference of overlapping communities, in: F. Pereira, C. J. C. Burges, L. Bottou, K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems*, volume 25, Neural Information Processing Systems Foundation, Inc., 2012, pp. 2249–2257.

- [498] D. I. Kim, P. K. Gopalan, D. Blei, E. Sudderth, Efficient online inference for Bayesian nonparametric relational models, in: *Advances in Neural Information Processing Systems*, volume 26, Neural Information Processing Systems Foundation, Inc., 2013, pp. 962–970.
- [499] K. Hayashi, T. Konishi, T. Kawamoto, A tractable fully Bayesian method for the stochastic block model, 2016. e-print arXiv:1602.02256.
- [500] C. Matias, T. Rebafka, F. Villers, A semiparametric extension of the stochastic block model for longitudinal networks, *Biometrika* 105 (2018) 665–680.
- [501] B. Perozzi, R. Al-Rfou, S. Skiena, DeepWalk: Online learning of social representations, in: S. Macskassy, C. Perlich, J. Leskovec, W. Wang, R. Ghani (Eds.), *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, Association for Computing Machinery, 2014, p. 701–710.
- [502] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, S. W. Zucker, Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps, *Proc. Natl. Acad. Sci. USA* 102 (2005) 7426–7431.
- [503] A. Grover, J. Leskovec, node2vec: Scalable feature learning for networks, in: B. Krishnapuram, M. Shah, A. Smola, C. Aggarwal, D. Shen, R. Rastogi (Eds.), *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, Association for Computing Machinery, 2016, pp. 855–864.
- [504] A. Singhal, Introducing the Knowledge Graph: things, not strings, 2012. Google Blog. Available at: <https://blog.google/products/search/introducing-knowledge-graph-things-not/>.
- [505] J. Ugander, B. Karrer, L. Backstrom, C. Marlow, The anatomy of the Facebook social graph, 2011. e-print arXiv:1111.4503.
- [506] D. Madey, LinkedIn’s Vision for an Economic Graph: A Conversation With Jeff Weiner and Thomas Friedman, 2020. The Economic Graph blog. Available at: <http://blog.linkedin.com/2014/06/30/linke>

dins-vision-for-an-economic-graph-a-conversation-with-jeff-weiner-and-thomas-friedman/.

- [507] T. Walle, Introducing: The Real World Graph, 2020. Unacast blog. Available at: <https://www.unacast.com/post/introducing-the-real-world-graph>.
- [508] J. Milinovich, Introducing the Pinterest Taste Graph and enhanced targeting, 2017. Pinterest blog. Available at: <https://business.pinterest.com/en/blog/introducing-the-pinterest-taste-graph-and-enhanced-targeting>.
- [509] N. Sharma, The Origin of Data Information Knowledge Wisdom (DIKW) Hierarchy, 2008. Google Inc. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [510] M. C. Ballandies, M. M. Dapp, E. Pournaras, Decrypting distributed ledger design-taxonomy, classification and blockchain community evaluation, 2018. e-print arXiv:1811.03419.
- [511] X. Wang, X. Zha, W. Ni, R. P. Liu, Y. J. Guo, X. Niu, K. Zheng, Survey on blockchain for Internet of Things, *Comput. Commun.* 136 (2019) 10–29.
- [512] Y. Zhou, F. R. Yu, J. Chen, Y. Kuo, Cyber-Physical-Social Systems: A State-of-the-Art Survey, Challenges and Opportunities, *IEEE Commun. Surv. Tutor.* 22 (2019) 389–425.
- [513] M. Pticek, V. Podobnik, G. Jezic, Beyond the internet of things: the social networking of machines, *Int. J. Distrib. Sens. Netw.* 12 (2016) 8178417.
- [514] K. Skala, D. Davidovic, E. Afgan, I. Sovic, Z. Sojat, Scalable distributed computing hierarchy: Cloud, fog and dew computing, *Open J. Cloud Comput. (OJCC)* 2 (2015) 16–24.
- [515] E. Guberović, T. Lipić, I. Čavrak, Dew Intelligence: Federated learning perspective, in: H. Shahriar (Ed.), 2021 IEEE 45th Annual Computers, Software, and Applications Conference (COMPSAC), IEEE, 2021, pp. 1819–1824.

- [516] X. Wang, Y. Han, V. C. Leung, D. Niyato, X. Yan, X. Chen, Convergence of edge computing and deep learning: A comprehensive survey, *IEEE Commun. Surv. Tutor.* 22 (2020) 869–904.
- [517] S. Deng, H. Zhao, W. Fang, J. Yin, S. Dustdar, A. Y. Zomaya, Edge intelligence: the confluence of edge computing and artificial intelligence, *IEEE Internet Things J.* 7 (2020) 7457–7469.
- [518] W. Y. B. Lim, N. C. Luong, D. T. Hoang, Y. Jiao, Y.-C. Liang, Q. Yang, D. Niyato, C. Miao, Federated learning in mobile edge networks: A comprehensive survey, *IEEE Commun. Surv. Tutor.* 22 (2020) 2031–2063.
- [519] F. Wang, M. Zhang, X. Wang, X. Ma, J. Liu, Deep Learning for Edge Computing Applications: A State-of-the-Art Survey, *IEEE Access* 8 (2020) 58322–58336.
- [520] M. Tsvetkova, T. Yasseri, E. T. Meyer, J. B. Pickering, V. Engen, P. Walland, M. Lüders, A. Følstad, G. Bravos, Understanding human-machine networks: a cross-disciplinary survey, *ACM Comput. Surv.* 50 (2017) 1–35.
- [521] I. Rahwan, M. Cebrian, N. Obradovich, J. Bongard, J.-F. Bonnefon, C. Breazeal, J. W. Crandall, N. A. Christakis, I. D. Couzin, M. O. Jackson, et al., Machine behaviour, *Nature* 568 (2019) 477–486.
- [522] D. Lazer, A. Pentland, L. Adamic, S. Aral, A.-L. Barabasi, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, et al., Computational social science., *Science* 323 (2009) 721–723.
- [523] P. Holme, F. Liljeros, Mechanistic models in computational social science, *Front. Phys.* 3 (2015) 78.
- [524] F. Schweitzer, Sociophysics, *Phys. Today* 71 (2018) 40–46.
- [525] K. Bhattacharya, K. Kaski, Social physics: uncovering human behaviour from communication, *Adv. Phys.* X 4 (2019) 1527723.
- [526] J. Blumenstock, G. Cadamuro, R. On, Predicting poverty and wealth from mobile phone metadata, *Science* 350 (2015) 1073–1076.

- [527] J. Zhang, D. Brackbill, S. Yang, J. Becker, N. Herbert, D. Centola, Support or competition? How online social networks increase physical activity: a randomized controlled trial, *Prev. Med. Rep.* 4 (2016) 453–458.
- [528] A. T. Jebb, L. Tay, E. Diener, S. Oishi, Happiness, income satiation and turning points around the world, *Nat. Hum. Behav.* 2 (2018) 33–38.
- [529] A. Bogomolov, B. Lepri, J. Staiano, E. Letouzé, N. Oliver, F. Pianesi, A. Pentland, Moves on the street: Classifying crime hotspots using aggregated anonymized data on people dynamics, *Big Data* 3 (2015) 148–158.
- [530] J. Ginsberg, M. H. Mohebbi, R. S. Patel, L. Brammer, M. S. Smolinski, L. Brilliant, Detecting influenza epidemics using search engine query data, *Nature* 457 (2009) 1012–1014.
- [531] A. Wesolowski, N. Eagle, A. J. Tatem, D. L. Smith, A. M. Noor, R. W. Snow, C. O. Buckee, Quantifying the impact of human mobility on malaria, *Science* 338 (2012) 267–270.
- [532] R. Wilson, E. zu Erbach-Schoenberg, M. Albert, D. Power, S. Tudge, M. Gonzalez, S. Guthrie, H. Chamberlain, C. Brooks, C. Hughes, et al., Rapid and near real-time assessments of population displacement using mobile phone data following disasters: the 2015 Nepal Earthquake, *PLOS Curr.* 8 (2016) ecur-rents.dis.d073fbece328e4c39087bc086d694b5c.
- [533] F. Ofii, P. Meier, M. Imran, C. Castillo, D. Tuia, N. Rey, J. Briant, P. Millet, F. Reinhard, M. Parkan, et al., Combining human computing and machine learning to make sense of big (aerial) data for disaster response, *Big Data* 4 (2016) 47–59.
- [534] W. R. Hobbs, M. Burke, N. A. Christakis, J. H. Fowler, Online social integration is associated with reduced mortality risk, *Proc. Natl. Acad. Sci. USA* 113 (2016) 12980–12984.
- [535] B. Podobnik, M. Jusup, D. Kovac, H. E. Stanley, Predicting the rise of EU right-wing populism in response to unbalanced immigration, *Complexity* 2017 (2017).

- [536] B. Lepri, J. Staiano, D. Sangokoya, E. Letouzé, N. Oliver, The tyranny of data? The bright and dark sides of data-driven decision-making for social good, in: T. Cerquitelli, D. Quercia, F. Pasquale (Eds.), *Transparent data mining for big and small data*, Springer, 2017, pp. 3–24.
- [537] R. Abebe, K. Goldner, Mechanism design for social good, *AI Matters* 4 (2018) 27–34.
- [538] Z. R. Shi, C. Wang, F. Fang, Artificial intelligence for social good: A survey, 2020. e-print arXiv:2001.01818.
- [539] R. Vinuesa, H. Azizpour, I. Leite, M. Balaam, V. Dignum, S. Domisch, A. Felländer, S. D. Langhans, M. Tegmark, F. F. Nerini, The role of artificial intelligence in achieving the Sustainable Development Goals, *Nat. Commun.* 11 (2020) 1–10.
- [540] Y.-A. De Montjoye, C. A. Hidalgo, M. Verleysen, V. D. Blondel, Unique in the crowd: The privacy bounds of human mobility, *Sci. Rep.* 3 (2013) 1376.
- [541] Y.-A. De Montjoye, L. Radaelli, V. K. Singh, A. S. Pentland, Unique in the shopping mall: On the reidentifiability of credit card metadata, *Science* 347 (2015) 536–539.
- [542] F. Pasquale, *The black box society*, Harvard University Press, 2015.
- [543] D. Ruths, J. Pfeffer, Social media for large studies of behavior, *Science* 346 (2014) 1063–1064.
- [544] N. Xie, G. Ras, M. van Gerven, D. Doran, Explainable deep learning: A field guide for the uninitiated, 2020. e-print arXiv:2004.14545.
- [545] L. Floridi, Establishing the rules for building trustworthy AI, *Nat. Mach. Intell.* 1 (2019) 261–262.
- [546] A. Theodorou, V. Dignum, Towards ethical and socio-legal governance in AI, *Nat. Mach. Intell.* 2 (2020) 10–12.
- [547] P. Ala-Pietilä, et al., Ethics guidelines for trustworthy ai, 2019. Report by High-Level Expert Group on Artificial Intelligence. Available at:

https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=60419. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.

- [548] L. Zdeborová, Machine learning: New tool in the box, *Nat. Phys.* 13 (2017) 420–421.
- [549] H. Shirado, N. A. Christakis, Locally noisy autonomous agents improve global human coordination in network experiments, *Nature* 545 (2017) 370–374.
- [550] J. W. Crandall, M. Oudah, F. Ishowo-Oloko, S. Abdallah, J.-F. Bonnefon, M. Cebrian, A. Shariff, M. A. Goodrich, I. Rahwan, et al., Cooperating with machines, *Nat. Commun.* 9 (2018) 233.
- [551] M. Mohammadi, A. Al-Fuqaha, S. Sorour, M. Guizani, Deep learning for IoT big data and streaming analytics: A survey, *IEEE Commun. Surv. Tutor.* 20 (2018) 2923–2960.
- [552] S. Zhang, L. Yao, A. Sun, Y. Tay, Deep learning based recommender system: A survey and new perspectives, *ACM Comput. Surv.* 52 (2019) 1–38.
- [553] M. Raghu, E. Schmidt, A survey of deep learning for scientific discovery, 2020. e-print arXiv:2003.11755.
- [554] P. Mehta, M. Bukov, C.-H. Wang, A. G. Day, C. Richardson, C. K. Fisher, D. J. Schwab, A high-bias, low-variance introduction to machine learning for physicists, *Phys. Rep.* 810 (2019) 1–124.
- [555] Y. Bengio, *Learning deep architectures for AI*, Now Publishers, 2009.
- [556] Y. Bengio, A. Courville, P. Vincent, Representation learning: A review and new perspectives, *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (2013) 1798–1828.
- [557] L. Deng, D. Yu, Deep learning: methods and applications, *Found. Trends Signal Process.* 7 (2014) 197–387.
- [558] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (2015) 436–444.

- [559] J. Schmidhuber, Deep learning in neural networks: An overview, *Neural Netw.* 61 (2015) 85–117.
- [560] I. Goodfellow, Y. Bengio, A. Courville, Y. Bengio, *Deep learning*, The MIT Press, 2016.
- [561] W. Liu, Z. Wang, X. Liu, N. Zeng, Y. Liu, F. E. Alsaadi, A survey of deep neural network architectures and their applications, *Neurocomputing* 234 (2017) 11–26.
- [562] J. Fan, C. Ma, Y. Zhong, A selective overview of deep learning, 2019. e-print arXiv:1904.05526.
- [563] J. Bullock, A. Luccioni, K. H. Pham, C. S. N. Lam, M. Luengo-Oroz, Mapping the landscape of artificial intelligence applications against COVID-19, 2020. e-print arXiv:2003.11336.
- [564] S. Barua, Understanding Coronanomics: The economic implications of the coronavirus (COVID-19) pandemic, 2020. Social Science Research Network (SSRN) Paper No. 3566477. Available at: <https://ssrn.com/abstract=3566477>.
- [565] S. R. Baker, N. Bloom, S. J. Davis, S. J. Terry, Covid-induced economic uncertainty, 2020. National Bureau of Economic Research (NBER) Working Paper No. 26983. Available at: <https://www.nber.org/papers/w26983>.
- [566] M. Nicola, Z. Alsafi, C. Sohrabi, A. Kerwan, A. Al-Jabir, C. Iosifidis, M. Agha, R. Agha, The socio-economic implications of the coronavirus and COVID-19 pandemic: a review, *Int. J. Surg.* 78 (2020) 185–193.
- [567] Y. Zhang, H. Lyu, Y. Liu, X. Zhang, Y. Wang, J. Luo, Monitoring depression trend on Twitter during the COVID-19 pandemic, 2020. e-print arXiv:2007.00228.
- [568] A. Oliver, A. Odena, C. A. Raffel, E. D. Cubuk, I. Goodfellow, Realistic evaluation of deep semi-supervised learning algorithms, in: S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), *Advances in neural information processing systems*, 2018, pp. 3235–3246.

- [569] Y. Ouali, C. Hudelot, M. Tami, An overview of deep semi-supervised learning, 2020. e-print arXiv:2006.05278.
- [570] D. P. Kingma, M. Welling, An introduction to variational autoencoders, 2019. e-print arXiv:1906.02691.
- [571] J. Gui, Z. Sun, Y. Wen, D. Tao, J. Ye, A review on generative adversarial networks: Algorithms, theory, and applications, 2020. e-print arXiv:2001.06937.
- [572] V. François-Lavet, P. Henderson, R. Islam, M. G. Bellemare, J. Pineau, An introduction to deep reinforcement learning, 2018. e-print arXiv:1811.12560.
- [573] Y. Li, Deep reinforcement learning: An overview, 2017. e-print arXiv:1701.07274.
- [574] K. Arulkumaran, M. P. Deisenroth, M. Brundage, A. A. Bharath, Deep reinforcement learning: A brief survey, *IEEE Signal Process. Mag.* 34 (2017) 26–38.
- [575] J. Achiam, Spinning up in deep RL, 2018. Available at: <https://spinningup.openai.com/en/latest/#>.
- [576] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, W. Zaremba, OpenAI Gym, 2016. e-print arXiv:1606.01540.
- [577] P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, Y. Wu, P. Zhokhov, OpenAI Baselines, 2017. Available at: <https://github.com/openai/baselines>.
- [578] A. Stooke, P. Abbeel, rlpyt: A research code base for deep reinforcement learning in pytorch, 2019. e-print arXiv:1909.01500.
- [579] P. Hernandez-Leal, B. Kartal, M. E. Taylor, A survey and critique of multiagent deep reinforcement learning, *Auton. Agent Multi Agent Syst.* 33 (2019) 750–797.
- [580] P. Hernandez-Leal, B. Kartal, M. E. Taylor, A very condensed survey and critique of multiagent deep reinforcement learning, in: A. El Fallah-Seghrouchni, G. Sukthankar, B. An, N. Yorke-Smith

(Eds.), Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems, 2020, pp. 2146–2148.

- [581] R. Roscher, B. Bohn, M. F. Duarte, J. Garcke, Explainable machine learning for scientific insights and discoveries, *IEEE Access* 8 (2020) 42200–42216.
- [582] R. Iten, T. Metger, H. Wilming, L. Del Rio, R. Renner, Discovering physical concepts with neural networks, *Phys. Rev. Lett.* 124 (2020) 010508.
- [583] L. Zdeborová, Understanding deep learning is also a job for physicists, *Nat. Phys.* 16 (2020) 602–604.
- [584] S.-M. Udrescu, M. Tegmark, AI Feynman: A physics-inspired method for symbolic regression, *Sci. Adv.* 6 (2020) eaay2631.
- [585] T. Wu, M. Tegmark, Toward an artificial intelligence physicist for unsupervised learning, *Phys. Rev. E* 100 (2019) 033311.
- [586] A. Choudhary, J. F. Lindner, E. G. Holliday, S. T. Miller, S. Sinha, W. L. Ditto, Physics enhanced neural networks predict order and chaos, 2019. e-print arXiv:1912.01958.
- [587] M. Koch-Janusz, Z. Ringel, Mutual information, neural networks and the renormalization group, *Nat. Phys.* 14 (2018) 578–582.
- [588] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* 91 (2019) 045002.
- [589] Y. Bahri, J. Kadmon, J. Pennington, S. S. Schoenholz, J. Sohl-Dickstein, S. Ganguli, Statistical mechanics of deep learning, *Annu. Rev. Condens. Matter Phys.* 11 (2020) 501–528.
- [590] H. Mhaskar, Q. Liao, T. Poggio, Learning functions: when is deep better than shallow, 2016. e-print arXiv:1603.00988.
- [591] H. W. Lin, M. Tegmark, D. Rolnick, Why does deep and cheap learning work so well?, *J. Stat. Phys.* 168 (2017) 1223–1247.

- [592] S.-B. Lin, Generalization and expressivity for deep nets, *IEEE Trans. Neural Netw. Learn. Syst.* 30 (2018) 1392–1406.
- [593] A. Decelle, V. Martin-Mayor, B. Seoane, Learning a local symmetry with neural networks, *Phys. Rev. E* 100 (2019) 050102.
- [594] J. Pennington, Y. Bahri, Geometry of neural network loss surfaces via random matrix theory, in: D. Precup, Y. W. Teh (Eds.), *Proceedings of the 34th International Conference on Machine Learning*, 2017, pp. 2798–2806.
- [595] A. M. Saxe, Y. Bansal, J. Dapello, M. Advani, A. Kolchinsky, B. D. Tracey, D. D. Cox, On the information bottleneck theory of deep learning, “*J. Stat. Mech. Theory Exp.* 2019 (2019) 124020.
- [596] Z. Goldfeld, Y. Polyanskiy, The information bottleneck problem and its applications in machine learning, *IEEE J. Sel. Areas Inform. Theory* 1 (2020) 19–38.
- [597] Z. Piran, R. Shwartz-Ziv, N. Tishby, The dual information bottleneck, 2020. e-print arXiv:2006.04641.
- [598] A. Achille, G. Paolini, S. Soatto, Where is the information in a deep neural network?, 2019. e-print arXiv:1905.12213.
- [599] H. Hafez-Kolahi, S. Kasaei, M. Soleymani-Baghshah, Do compressed representations generalize better?, 2019. e-print arXiv:1909.09706.
- [600] F. D’Angelo, L. Böttcher, Learning the Ising model with generative neural networks, *Phys. Rev. Res.* 2 (2020) 023266.
- [601] V. Boucher, *MONTREAL.AI academy: artificial intelligence 101 first world-class overview of AI for all*, 2020. VIP AI 101 cheat sheet. Available at: <https://www.montreal.ai/ai4all.pdf>. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [602] S. Russell, P. Norvig, *Artificial intelligence: a modern approach*, Prentice Hall, 2002.
- [603] M. I. Jordan, Artificial intelligence—the revolution hasn’t happened yet, *Harvard Data Sci. Rev.* 1 (2019).

- [604] J. R. Searle, Minds, brains, and programs, in: D. J. Levitin (Ed.), Foundations of cognitive psychology: core readings, The MIT Press, 1980, pp. 417–457.
- [605] R. Fjelland, Why general artificial intelligence will not be realized, *Humanit. Soc. Sci. Commun.* 7 (2020) 10.
- [606] S. Shalev-Shwartz, S. Ben-David, Understanding machine learning: From theory to algorithms, Cambridge University Press, 2014.
- [607] Y. S. Abu-Mostafa, M. Magdon-Ismail, H.-T. Lin, Learning from data: a short course, AMLBook, 2012.
- [608] P. Domingos, A few useful things to know about machine learning, *Commun. ACM* 55 (2012) 78–87.
- [609] K. P. Bennett, E. Parrado-Hernández, The interplay of optimization and machine learning research, *J. Mach. Learn. Res.* 7 (2006) 1265–1281.
- [610] A. Esteva, A. Robicquet, B. Ramsundar, V. Kuleshov, M. DePristo, K. Chou, C. Cui, G. Corrado, S. Thrun, J. Dean, A guide to deep learning in healthcare, *Nat. Med.* 25 (2019) 24–29.
- [611] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Comput.* 9 (1997) 1735–1780.
- [612] J. Chung, C. Gulcehre, K. Cho, Y. Bengio, Empirical evaluation of gated recurrent neural networks on sequence modeling, 2014. e-print arXiv:1412.3555.
- [613] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, 2017. e-print arXiv:1706.03762.
- [614] Y. Qin, D. Song, H. Chen, W. Cheng, G. Jiang, G. Cottrell, A dual-stage attention-based recurrent neural network for time series prediction, 2017. e-print arXiv:1704.02971.
- [615] I. R. Ward, J. Joyner, C. Lickfold, S. Rowe, Y. Guo, M. Ben-namoun, A practical guide to graph neural networks, 2020. e-print arXiv:2010.05234.

- [616] P. Stone, R. Brooks, E. Brynjolfsson, R. Calo, O. Etzioni, G. Hager, J. Hirschberg, S. Kalyanakrishnan, E. Kamar, S. Kraus, et al., Artificial intelligence and life in 2030: the one hundred year study on artificial intelligence, 2016. Available at: https://ai100.stanford.edu/sites/g/files/sbiybj9861/f/ai100report10032016fnl_singles.pdf.
- [617] E. Turkina, The importance of networking to entrepreneurship: Montreal’s artificial intelligence cluster and its born-global firm element ai, *J. Small Bus. Entrepreneurship* 30 (2018) 1–8.
- [618] E. Cetinic, T. Lipic, S. Grgic, Fine-tuning convolutional neural networks for fine art classification, *Expert Syst. Appl.* 114 (2018) 107–118.
- [619] J. Petangoda, N. A. Monk, M. P. Deisenroth, A foliated view of transfer learning, 2020. e-print arXiv:2008.00546.
- [620] P. W. Battaglia, R. Pascanu, M. Lai, D. Rezende, K. Kavukcuoglu, Interaction networks for learning about objects, relations and physics, 2016. e-print arXiv:1612.00222.
- [621] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, S. J. Gershman, Building machines that learn and think like people, *Behav. Brain Sci.* 40 (2017) e253.
- [622] B. M. Lake, R. Salakhutdinov, J. B. Tenenbaum, The Omniglot challenge: a 3-year progress report, *Curr. Opin. Behav. Sci.* 29 (2019) 97–104.
- [623] T. Hospedales, A. Antoniou, P. Micaelli, A. Storkey, Meta-learning in neural networks: A survey, 2020. e-print arXiv:2004.05439.
- [624] Y. Wang, Q. Yao, J. T. Kwok, L. M. Ni, Generalizing from a few examples: A survey on few-shot learning, *ACM Comput. Surv.* 53 (2020) 1–34.
- [625] G. Koch, Siamese neural networks for one-shot image recognition, 2015. Master thesis, University of Toronto. Available at: <http://www.cs.toronto.edu/~gkoch/files/msc-thesis.pdf>.

- [626] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, D. Wierstra, Matching networks for one shot learning, e-print arXiv:1606.04080 (2016).
- [627] J. Snell, K. Swersky, R. S. Zemel, Prototypical networks for few-shot learning, e-print arXiv:1703.05175 (2017).
- [628] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, T. M. Hospedales, Learning to compare: Relation network for few-shot learning, in: D. Forsyth, I. Laptev, A. Oliva, D. Ramanan, M. S. Brown, B. Morse, S. Peleg (Eds.), 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), IEEE Computer Society, 2018, pp. 1199–1208.
- [629] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, T. Lillicrap, Meta-learning with memory-augmented neural networks, in: M. F. Balcan, K. Q. Weinberger (Eds.), Proceedings of the 33rd International Conference on Machine Learning, volume 48, Proceedings of Machine Learning Research, 2016, pp. 1842–1850.
- [630] S. Ravi, H. Larochelle, Optimization as a model for few-shot learning, 2016. International Conference on Learning Representations 2017. Available at: <https://openreview.net/pdf?id=rJY0-Kc11>. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [631] C. Finn, P. Abbeel, S. Levine, Model-agnostic meta-learning for fast adaptation of deep networks, 2017. e-print arXiv:1703.03400.
- [632] A. Raghu, M. Raghu, S. Bengio, O. Vinyals, Rapid learning or feature reuse? Towards understanding the effectiveness of MAML, 2019. e-print arXiv:1909.09157.
- [633] X. Liang, Z. Yan, A survey on game theoretical methods in human-machine networks, *Future Gener. Comput. Syst.* 92 (2019) 674–693.
- [634] J. N. Foerster, R. Y. Chen, M. Al-Shedivat, S. Whiteson, P. Abbeel, I. Mordatch, Learning with opponent-learning awareness, 2017. e-print arXiv:1709.04326.

- [635] W. H. Press, F. J. Dyson, Iterated Prisoner’s Dilemma contains strategies that dominate any evolutionary opponent, *Proc. Natl. Acad. Sci. USA* 109 (2012) 10409–10413.
- [636] A. J. Stewart, J. B. Plotkin, Extortion and cooperation in the Prisoner’s Dilemma, *Proc. Natl. Acad. Sci. USA* 109 (2012) 10134–10135.
- [637] N. Rabinowitz, F. Perbet, F. Song, C. Zhang, S. M. A. Eslami, M. Botvinick, Machine theory of mind, in: J. Dy, A. Krause (Eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80, *Proceedings of Machine Learning Research*, 2018, pp. 4218–4227.
- [638] N. E. Glynatsi, V. A. Knight, Using a theory of mind to find best responses to memory-one strategies, *Sci. Rep.* 10 (2020) 17287.
- [639] V. A. Knight, O. Campbell, M. Harper, K. M. Langner, J. Campbell, T. Campbell, A. Carney, M. Chorley, C. Davidson-Pilon, K. Glass, et al., An open framework for the reproducible study of the iterated prisoner’s dilemma, *J. Open Res. Softw.* 4 (2016) e35.
- [640] M. Harper, V. Knight, M. Jones, G. Koutsovoulos, N. E. Glynatsi, O. Campbell, Reinforcement learning produces dominant strategies for the Iterated Prisoner’s Dilemma, *PLOS ONE* 12 (2017) e0188046.
- [641] T.-C. Au, D. Nau, Accident or intention: that is the question (in the Noisy Iterated Prisoner’s Dilemma), in: H. Nakashima, M. Wellman, G. Weiss, P. Stone (Eds.), *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, Association for Computing Machinery, 2006, pp. 561–568.
- [642] J. Z. Leibo, V. Zambaldi, M. Lanctot, J. Marecki, T. Graepel, Multi-agent reinforcement learning in sequential social dilemmas, e-print arXiv:1702.03037 (2017).
- [643] A. Peysakhovich, A. Lerer, Towards AI that can solve social dilemmas, in: C. Amato, T. Graepel, J. Leibo, F. Oliehoek, K. Tuyls (Eds.), *The 2018 AAAI Spring Symposium Series*, volume SS-18, Association for the Advancement of Artificial Intelligence, 2018, pp. 649–654.

- [644] R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel, I. Mordatch, Multi-agent actor-critic for mixed cooperative-competitive environments, e-print arXiv:1706.02275 (2017).
- [645] Y. J. Park, Y. S. Cho, S. B. Kim, Multi-agent reinforcement learning with approximate model learning for competitive games, PLOS ONE 14 (2019) e0222215.
- [646] Y. Bachrach, R. Everett, E. Hughes, A. Lazaridou, J. Z. Leibo, M. Lanctot, M. Johanson, W. M. Czarnecki, T. Graepel, Negotiating team formation using deep reinforcement learning, Artif. Intell. 288 (2020) 103356.
- [647] T. Degris, M. White, R. S. Sutton, Off-policy actor-critic, e-print arXiv:1205.4839 (2012).
- [648] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, et al., Soft actor-critic algorithms and applications, e-print arXiv:1812.05905 (2018).
- [649] V. R. Konda, J. N. Tsitsiklis, Actor-critic algorithms, in: S. A. Solla, T. K. Leen, K.-R. Müller (Eds.), Advances in neural information processing systems, volume 12, The MIT Press, 2000, pp. 1008–1014.
- [650] L. Kraemer, B. Banerjee, Multi-agent reinforcement learning as a rehearsal for decentralized planning, Neurocomputing 190 (2016) 82–94.
- [651] A. OroojlooyJadid, D. Hajinezhad, A review of cooperative multi-agent deep reinforcement learning, e-print arXiv:1908.03963 (2019).
- [652] K. Zhang, Z. Yang, T. Başar, Multi-agent reinforcement learning: A selective overview of theories and algorithms, e-print arXiv:1911.10635 (2019).
- [653] W. Li, B. Jin, X. Wang, J. Yan, H. Zha, F2A2: Flexible Fully-decentralized Approximate Actor-critic for cooperative multi-agent reinforcement learning, e-print arXiv:2004.11145 (2020).
- [654] E. A. Hansen, D. S. Bernstein, S. Zilberstein, Dynamic programming for partially observable stochastic games, in: A. G. Cohn (Ed.),

Proceedings of the 19th national conference on Artificial intelligence, volume 4, AAAI Press, 2004, pp. 709–715.

- [655] M. Rovatsos, We may not cooperate with friendly machines, *Nat. Mach. Intell.* 1 (2019) 497–498.
- [656] F. Ishowo-Oloko, J.-F. Bonnefon, Z. Soroye, J. Crandall, I. Rahwan, T. Rahwan, Behavioural evidence for a transparency-efficiency trade-off in human-machine cooperation, *Nat. Mach. Intell.* 1 (2019) 517–521.
- [657] E. Yoeli, M. Hoffman, D. G. Rand, M. A. Nowak, Powering up with indirect reciprocity in a large-scale field experiment, *Proc. Natl. Acad. Sci. USA* 110 (2013) 10424–10429.
- [658] A. Lerer, A. Peysakhovich, Maintaining cooperation in complex social dilemmas using deep reinforcement learning, e-print arXiv:1707.01068 (2017).
- [659] J.-F. Bonnefon, I. Rahwan, Machine thinking, fast and slow, *Trends Cogn. Sci.* 24 (2020) 1019–1027.
- [660] R. Karandikar, D. Mookherjee, D. Ray, F. Vega-Redondo, Evolving aspirations and cooperation, *J. Econ. Theory* 80 (1998) 292–331.
- [661] J. Stoyanovich, J. J. Van Bavel, T. V. West, The imperative of interpretable machines, *Nat. Mach. Intell.* 2 (2020) 197–199.
- [662] M. Kearns, S. Suri, N. Montfort, An experimental study of the coloring problem on human subject networks, *Science* 313 (2006) 824–827.
- [663] E. Ferrara, Disinformation and social bot operations in the run up to the 2017 French presidential election, e-print arXiv:1707.00086 (2017).
- [664] D. M. Lazer, M. A. Baum, Y. Benkler, A. J. Berinsky, K. M. Greenhill, F. Menczer, M. J. Metzger, B. Nyhan, G. Pennycook, D. Rothschild, et al., The science of fake news, *Science* 359 (2018) 1094–1096.
- [665] M. Stella, E. Ferrara, M. De Domenico, Bots increase exposure to negative and inflammatory content in online social systems, *Proc. Natl. Acad. Sci. USA* 115 (2018) 12435–12440.

- [666] A. J. Stewart, M. Mosleh, M. Diakonova, A. A. Arechar, D. G. Rand, J. B. Plotkin, Information gerrymandering and undemocratic decisions, *Nature* 573 (2019) 117–121.
- [667] K. Crawford, R. Calo, There is a blind spot in AI research, *Nature* 538 (2016) 311–313.
- [668] M. R. D’Orsogna, M. Perc, Statistical physics of crime: A review, *Phys. Life Rev.* 12 (2015) 1–21.
- [669] P. Ball, *Why society is a complex matter*, Springer, 2012.
- [670] D. Helbing, D. Brockmann, T. Chadeaux, K. Donnay, U. Blanke, O. Woolley-Meza, M. Moussaid, A. Johansson, J. Krause, S. Schutte, et al., Saving human lives: What complexity science and information systems can contribute, *J. Stat. Phys.* 158 (2015) 735–781.
- [671] J. Q. Wilson, G. L. Kelling, Broken windows, *Atlantic Monthly* 249 (1982) 29–38.
- [672] Y. A. Kuznetsov, *Elements of applied bifurcation theory*, Springer, 2004.
- [673] H. E. Stanley, *Introduction to Phase Transitions and Critical Phenomena*, Oxford University Press, 1971.
- [674] C. Castellano, S. Fortunato, V. Loreto, Statistical physics of social dynamics, *Rev. Mod. Phys.* 81 (2009) 591–646.
- [675] L. E. Cohen, M. Felson, Social change and crime rate trends: A routine activity approach, *Am. Sociol. Rev.* (1979) 588–608.
- [676] M. B. Short, M. R. D’orsogna, V. B. Pasour, G. E. Tita, P. J. Brantingham, A. L. Bertozzi, L. B. Chayes, A statistical model of criminal behavior, *Math. Mod. Meth. Appl. Sci.* 18 (2008) 1249–1267.
- [677] S. D. Johnson, K. Bowers, A. Hirschfield, New insights into the spatial and temporal distribution of repeat victimization, *Brit. J. Criminol.* 37 (1997) 224–241.

- [678] M. Townsley, R. Homel, J. Chaseling, Repeat burglary victimisation: Spatial and temporal patterns, *Austr. New. Zeal. J. Criminol.* 33 (2000) 37–63.
- [679] S. D. Johnson, K. J. Bowers, The stability of space-time clusters of burglary, *Brit. J. Criminol.* 44 (2004) 55–65.
- [680] M. B. Short, M. R. D’orsogna, P. J. Brantingham, G. E. Tita, Measuring and modeling repeat and near-repeat burglary effects, *J. Quant. Criminol.* 25 (2009) 325–339.
- [681] L. Green, Cleaning up drug hot spots in Oakland, California: The displacement and diffusion effects, *Just. Quart.* 12 (1995) 737–754.
- [682] A. A. Braga, The effects of hot spots policing on crime, *Ann. Am. Acad. Pol. Soc. Sci.* 578 (2001) 104–125.
- [683] A. A. Braga, A. V. Papachristos, D. M. Hureau, The effects of hot spots policing on crime: An updated systematic review and meta-analysis, *Just. Quart.* 31 (2014) 633–663.
- [684] D. Weisburd, L. A. Wyckoff, J. Ready, J. E. Eck, J. C. Hinkle, F. Gajewski, Does crime just move around the corner? A controlled study of spatial displacement and diffusion of crime control benefits, *Criminology* 44 (2006) 549–592.
- [685] T. A. Taniguchi, G. F. Rengert, E. S. McCord, Where size matters: Agglomeration economies of illegal drug markets in Philadelphia, *Just. Quart.* 26 (2009) 670–694.
- [686] G. S. Becker, Crime and punishment: An economic approach, *J. Political Econ.* 76 (1968) 169–217.
- [687] A. N. Doob, C. M. Webster, Sentence severity and crime: Accepting the null hypothesis, *Crime Just.* 30 (2003) 143–195.
- [688] S. D. Johnson, W. Bernasco, K. J. Bowers, H. Elffers, J. Ratcliffe, G. Rengert, M. Townsley, Space-time patterns of risk: a cross national assessment of residential burglary victimization, *J. Quant. Criminol.* 23 (2007) 201–219.

- [689] G. F. Rengert, A. R. Piquero, P. R. Jones, Distance decay reexamined, *Criminology* 37 (1999) 427–446.
- [690] W. Bernasco, F. Luykx, Effects of attractiveness, opportunity and accessibility to burglars on residential burglary rates of urban neighborhoods, *Criminology* 41 (2003) 981–1002.
- [691] W. Bernasco, P. Nieuwbeerta, How do residential burglars select target areas? a new approach to the analysis of criminal location choice, *Brit. J. Criminol.* 45 (2005) 296–315.
- [692] M. B. Short, P. J. Brantingham, A. L. Bertozzi, G. E. Tita, Dissipation and displacement of hotspots in reaction-diffusion models of crime, *Proc. Natl. Acad. Sci. USA* 107 (2010) 3961–3965.
- [693] M. B. Short, A. L. Bertozzi, P. J. Brantingham, Nonlinear patterns in urban crime: Hotspots, bifurcations, and suppression, *SIAM J. Appl. Dyn. Syst.* 9 (2010) 462–483.
- [694] M. C. Cross, P. C. Hohenberg, Pattern formation outside of equilibrium, *Rev. Mod. Phys.* 65 (1993) 851.
- [695] N. Rodriguez, A. Bertozzi, Local existence and uniqueness of solutions to a PDE model for criminal behavior, *Math. Mod. Meth. Appl. Sci.* 20 (2010) 1425–1457.
- [696] P. A. Jones, P. J. Brantingham, L. R. Chayes, Statistical models of criminal behavior: the effects of law enforcement actions, *Math. Mod. Meth. Appl. Sci.* 20 (2010) 1397–1423.
- [697] R. S. Cantrell, C. Cosner, R. Manásevich, Global bifurcation of solutions for crime modeling equations, *SIAM J. Math. Anal.* 44 (2012) 1340–1358.
- [698] H. Berestycki, N. Rodriguez, L. Ryzhik, Traveling wave solutions in a reaction-diffusion model for criminal activity, *Multiscale Model. Simul.* 11 (2013) 1097–1126.
- [699] J. R. Zipkin, M. B. Short, A. L. Bertozzi, Cops on the dots in a mathematical model of urban crime and police response, *Disc. Cont. Dyn. Syst. B* 19 (2014) 1479–1506.

- [700] J. C. Nuno, M. A. Herrero, M. Primicerio, A triangle model of criminality, *Physica A* 387 (2008) 2926–2936.
- [701] J. C. Nuno, M. A. Herrero, M. Primicerio, A mathematical model of a criminal-prone society, *Disc. Cont. Dyn. Syst. S* 4 (2011) 193–207.
- [702] D. McMillon, C. P. Simon, J. Morenoff, Modeling the underlying dynamics of the spread of crime, *PLOS ONE* 9 (2014) e88923.
- [703] J. Sooknanan, B. Bhatt, D. M. G. Comissiong, Criminals treated as predators to be harvested: a two prey one predator model with group defense, prey migration and switching, *J. Math. Res.* 4 (2012) 92–106.
- [704] G. O. Mohler, M. B. Short, P. J. Brantingham, F. P. Schoenberg, G. E. Tita, Self-exciting point process modeling of crime, *J. Am. Stat. Assoc.* 106 (2011) 100–108.
- [705] D. J. Daley, D. Vere-Jones, An introduction to the theory of point processes, volume 1: Elementary theory and methods, Springer, 2003.
- [706] K. J. Bowers, S. D. Johnson, K. Pease, Prospective hot-spotting: the future of crime mapping?, *Brit. J. Criminol.* 44 (2004) 641–658.
- [707] E. Lewis, G. Mohler, P. J. Brantingham, A. L. Bertozzi, Self-exciting point process models of civilian deaths in Iraq, *Secur. J.* 25 (2012) 244–264.
- [708] G. O. Mohler, M. B. Short, Geographic profiling from kinetic models of criminal behavior, *SIAM J. Appl. Math.* 72 (2012) 163–180.
- [709] R. Axelrod, *The Evolution of Cooperation*, Basic Books, 1984.
- [710] G. Hardin, The tragedy of the commons, *Science* 162 (1968) 1243–1248.
- [711] M. B. Short, P. J. Brantingham, M. R. D’orsogna, Cooperation and punishment in an adversarial game: How defectors pave the way to a peaceful society, *Phys. Rev. E* 82 (2010) 066114.
- [712] J. P. Lynch, L. A. Addington, *Understanding Crime Statistics: Revisiting the Divergence of the NCVS and the UCR*, Cambridge University Press, 2006.

- [713] E. Fehr, Don't lose your reputation, *Nature* 432 (2004) 449–450.
- [714] A. Szolnoki, M. Perc, Second-order free-riding on antisocial punishment restores the effectiveness of prosocial punishment, *Phys. Rev. X* 7 (2017) 041027.
- [715] K. H. Schlag, Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits, *J. Econ. Theory* 78 (1998) 130–156.
- [716] M. B. Short, A. B. Pitcher, M. R. D'orsogna, External conversions of player strategy in an evolutionary game: A cost-benefit analysis through optimal control, *Eur. J. Appl. Math.* 24 (2013) 131.
- [717] M. R. D'Orsogna, R. Kendall, M. McBride, M. B. Short, Criminal defectors lead to the emergence of cooperation in an experimental, adversarial game, *PLOS ONE* 8 (2013) e61458.
- [718] G. Tsebelis, Penalty has no impact on crime: A game-theoretic analysis, *Ration. Soc.* 2 (1990) 255–286.
- [719] M. Perc, K. Donnay, D. Helbing, Understanding recurrent crime as system-immanent collective behavior, *PLOS ONE* 8 (2013) e76063.
- [720] M. Perc, A. Szolnoki, A double-edged sword: Benefits and pitfalls of heterogeneous punishment in evolutionary inspection games, *Sci. Rep.* 5 (2015) 11027.
- [721] D. McFadden, Conditional logit analysis of qualitative choice behavior, in: P. Zarembka (Ed.), *Frontiers in econometrics*, Academic Press, 1974, pp. 105–142.
- [722] G. Szabó, G. Fath, Evolutionary games on graphs, *Phys. Rep.* 446 (2007) 97–216.
- [723] M. Perc, Stability of subsystem solutions in agent-based models, *Eur. J. Phys.* 39 (2017) 014001.
- [724] S. N. Dorogovtsev, J. F. Mendes, Evolution of networks, *Adv. Phys.* 51 (2002) 1079–1187.

- [725] E. Estrada, *The Structure of Complex Networks: Theory and Applications*, Oxford University Press, 2012.
- [726] S. L. Mallory, *Understanding organized crime*, Jones & Bartlett Publishers, 2011.
- [727] H. V. Ribeiro, L. G. A. Alves, A. F. Martins, E. K. Lenzi, M. Perc, The dynamical structure of political corruption networks, *J. Complex Netw.* 6 (2018) 989–1003.
- [728] D. Gambetta, *Trust: Making and breaking cooperative relations*, B. Blackwell, 1988.
- [729] J. S. Beittel, Mexico’s drug-related violence, 2009. Report by Congressional Research Service. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [730] S. Raphael, R. Winter-Ebmer, Identifying the effect of unemployment on crime, *J. Law Econ.* 44 (2001) 259–283.
- [731] M.-J. Lin, Does unemployment increase crime? Evidence from US data 1974–2000, *J. Hum. Resour.* 43 (2008) 413–436.
- [732] G. LaFree, Declining violent crime rates in the 1990s: Predicting crime booms and busts, *Ann. Rev. Sociol.* 25 (1999) 145–168.
- [733] R. Curtis, The improbable transformation of inner-city neighborhoods: Crime, violence, drugs, and youth in the 1990s, *J. Crim. Law Criminol.* 88 (1997) 1233.
- [734] G. LaFree, *Losing legitimacy: Street crime and the decline of social institutions in America*, Westview Press, 1998.
- [735] N. Corsaro, E. F. McGarrell, Testing a promising homicide reduction strategy: Re-assessing the impact of the Indianapolis “pulling levers” intervention, *J. Exp. Criminol.* 5 (2009) 63.
- [736] E. F. McGarrell, N. Corsaro, N. K. Hipple, T. S. Bynum, Project safe neighborhoods and violent crime trends in US cities: Assessing violent crime impact, *J. Quant. Criminol.* 26 (2010) 165–190.

- [737] F. Calderoni, G. M. Campedelli, A. Szekely, M. Paolucci, G. Andrighetto, Recruitment into organized crime: An agent-based approach testing the impact of different policies, *J. Quant. Criminol.* (2021).
- [738] G. A. Gómez-Sorzano, Decomposing violence: crime cycles in the twentieth century in the United States, *Appl. Econ. Int. Dev.* 7 (2007).
- [739] F. E. Zimring, *The great American crime decline*, Oxford University Press, 2006.
- [740] A. V. Papachristos, A. A. Braga, D. M. Hureau, Social networks and the risk of gunshot injury, *J. Urban Health* 89 (2012) 992–1003.
- [741] A. V. Papachristos, D. M. Hureau, A. A. Braga, The corner and the crew: The influence of geography and social networks on gang violence, *Am. Sociol. Rev.* 78 (2013) 417–447.
- [742] R. Albert, H. Jeong, A.-L. Barabási, Error and attack tolerance of complex networks, *Nature* 406 (2000) 378–382.
- [743] P. A. Duijn, V. Kashirin, P. M. Sloot, The relative ineffectiveness of criminal network disruption, *Sci. Rep.* 4 (2014) 4238.
- [744] R. A. Hegemann, L. M. Smith, A. B. Barbaro, A. L. Bertozzi, S. E. Reid, G. E. Tita, Geographical influences of an emerging network of gang rivalries, *Physica A* 390 (2011) 3894–3914.
- [745] Y. van Gennip, B. Hunter, R. Ahn, P. Elliott, K. Luh, M. Halvorson, S. Reid, M. Valasik, J. Wo, G. E. Tita, et al., Community detection using spectral clustering on sparse geosocial data, *SIAM J. Appl. Math.* 73 (2013) 67–83.
- [746] A. B. Barbaro, L. Chayes, M. R. D’Orsogna, Territorial developments based on graffiti: A statistical mechanics approach, *Physica A* 392 (2013) 252–270.
- [747] K. Thompson, N. Offler, L. Hirsch, D. Every, M. J. Thomas, D. Dawson, From broken windows to a renovated research agenda: A review of the literature on vandalism and graffiti in the rail industry, *Transp. Res. Part A Policy Pract.* 46 (2012) 1280–1290.

- [748] S. Rose-Ackerman, The economics of corruption, *J. Public Econ.* 4 (1975) 187–203.
- [749] A. Shleifer, R. W. Vishny, Corruption, *Q. J. Econ.* 108 (1993) 599–617.
- [750] P. Mauro, Corruption and growth, *Q. J. Econ.* 110 (1995) 681–712.
- [751] P. Bardhan, Corruption and development: a review of issues, *J. Econ. Lit.* 35 (1997) 1320–1346.
- [752] J. Shao, P. C. Ivanov, B. Podobnik, H. E. Stanley, Quantitative relations between corruption and economic factors, *Eur. Phys. J. B* 56 (2007) 157–166.
- [753] M. E. Haque, R. Kneller, et al., Public investment and growth: The role of corruption, 2008. Centre for Growth and Business Cycle Research discussion paper series 098. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [754] S. Gupta, H. Davoodi, R. Alonso-Terme, Does corruption affect income inequality and poverty?, *Econ. Gov.* 3 (2002) 23–45.
- [755] J. Xu, H. Chen, Criminal network analysis and visualization, *Commun. ACM* 48 (2005) 100–107.
- [756] M. Del Vicario, A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H. E. Stanley, W. Quattrociocchi, The spreading of misinformation online, *Proc. Natl. Acad. Sci. USA* 113 (2016) 554–559.
- [757] I. Luna-Pla, J. R. Nicolás-Carlock, Corruption and complexity: a scientific framework for the analysis of corruption networks, *Appl. Netw. Sci* 5 (2020) 13.
- [758] J. R. Nicolás-Carlock, I. Luna-Pla, Conspiracy of corporate networks in corruption scandals, *Front. Phys.* 9 (2021) 301.
- [759] R. Guimera, L. A. N. Amaral, Cartography of complex networks: modules and universal roles, *J. Stat. Mech. Theory Exp.* 2005 (2005) P02001.

- [760] E. Ferrara, P. De Meo, S. Catanese, G. Fiumara, Detecting criminal organizations in mobile phone networks, *Expert Syst. Appl.* 41 (2014) 5733–5750.
- [761] A. Bogomolov, B. Lepri, J. Staiano, N. Oliver, F. Pianesi, A. Pentland, Once upon a crime: towards crime prediction from demographics and mobile data, in: A. A. Salah, J. Cohn, B. Schuller, O. Aran, L.-P. Morency, P. R. Cohen (Eds.), *Proceedings of the 16th international conference on multimodal interaction*, Association for Computing Machinery, 2014, pp. 427–434.
- [762] B. Herrmann, C. Thöni, S. Gächter, Antisocial punishment across societies, *Science* 319 (2008) 1362–1367.
- [763] D. G. Rand, A. Dreber, T. Ellingsen, D. Fudenberg, M. A. Nowak, Positive interactions promote public cooperation, *Science* 325 (2009) 1272–1275.
- [764] M. Milinski, D. Semmann, H.-J. Krambeck, Reputation helps solve the ‘tragedy of the commons’, *Nature* 415 (2002) 424–426.
- [765] D. G. Rand, M. A. Nowak, The evolution of antisocial punishment in optional public goods games, *Nat. Commun.* 2 (2011) 434.
- [766] S. Gächter, E. Renner, M. Sefton, The long-run benefits of punishment, *Science* 322 (2008) 1510–1510.
- [767] R. Boyd, H. Gintis, S. Bowles, Coordinated punishment of defectors sustains cooperation and can proliferate when rare, *Science* 328 (2010) 617–620.
- [768] M. Perc, A. Szolnoki, Self-organization of punishment in structured populations, *New J. Phys.* 14 (2012) 043013.
- [769] C. Hilbe, K. Sigmund, Incentives and opportunism: from the carrot to the stick, *Proc. R. Soc. B* 277 (2010) 2427–2433.
- [770] A. Szolnoki, M. Perc, Reward and cooperation in the spatial public goods game, *EPL (Europhys. Lett.)* 92 (2010) 38003.
- [771] A. Szolnoki, M. Perc, Evolutionary advantages of adaptive rewarding, *New J. Physics* 14 (2012) 093016.

- [772] A. Szolnoki, M. Perc, Antisocial pool rewarding does not deter public cooperation, *Proc. R. Soc. B* 282 (2015) 20151975.
- [773] Y. Fang, T. P. Benko, M. Perc, H. Xu, Q. Tan, Synergistic third-party rewarding and punishment in the public goods game, *Proc. R. Soc. A* 475 (2019) 20190349.
- [774] L. Hu, N. He, Q. Weng, X. Chen, M. Perc, Rewarding endowments lead to a win-win in the evolution of public cooperation and the accumulation of common resources, *Chaos Solitons Fractals* 134 (2020) 109694.
- [775] A. Szolnoki, M. Perc, Correlation of positive and negative reciprocity fails to confer an evolutionary advantage: Phase transitions to elementary strategies, *Phys. Rev. X* 3 (2013) 041021.
- [776] K. Sigmund, H. De Silva, A. Traulsen, C. Hauert, Social learning promotes institutions for governing the commons, *Nature* 466 (2010) 861–863.
- [777] A. Szolnoki, G. Szabó, M. Perc, Phase diagrams for the spatial public goods game with pool punishment, *Phys. Rev. E* 83 (2011) 036101.
- [778] M. Perc, Sustainable institutionalized punishment requires elimination of second-order free-riders, *Sci. Rep.* 2 (2012) 1–6.
- [779] X. Chen, M. Perc, Optimal distribution of incentives for public cooperation in heterogeneous interaction environments, *Front. Behav. Neurosci.* 8 (2014) 248.
- [780] B. Berenji, T. Chou, M. R. D’Orsogna, Recidivism and rehabilitation of criminal offenders: A carrot and stick evolutionary game, *PLOS ONE* 9 (2014) e85531.
- [781] S. N. Dorogovtsev, A. V. Goltsev, J. F. Mendes, Critical phenomena in complex networks, *Rev. Mod. Phys.* 80 (2008) 1275.
- [782] M. Scheffer, S. R. Carpenter, T. M. Lenton, J. Bascompte, W. Brock, V. Dakos, J. Van de Koppel, I. A. Van de Leemput, S. A. Levin, E. H. Van Nes, et al., Anticipating critical transitions, *Science* 338 (2012) 344–348.

- [783] Y. Yu, G. Xiao, J. Zhou, Y. Wang, Z. Wang, J. Kurths, H. J. Schellnhuber, System crash as dynamics of complex networks, *Proc. Natl. Acad. Sci. USA* 113 (2016) 11726–11731.
- [784] R. Cohen, K. Erez, D. Ben-Avraham, S. Havlin, Resilience of the internet to random breakdowns, *Phys. Rev. Lett.* 85 (2000) 4626.
- [785] A. Majdandzic, L. A. Braunstein, C. Curme, I. Vodenska, S. Levy-Carciente, H. E. Stanley, S. Havlin, Multiple tipping points and optimal repairing in interacting networks, *Nat. Commun.* 7 (2016) 10850.
- [786] A. Majdandzic, B. Podobnik, S. V. Buldyrev, D. Y. Kenett, S. Havlin, H. E. Stanley, Spontaneous recovery in dynamical networks, *Nat. Phys.* 10 (2014) 34–38.
- [787] K. M. Greenhill, Open arms behind barred doors: fear, hypocrisy and policy schizophrenia in the European migration crisis, *Eur. Law J.* 22 (2016) 317–332.
- [788] S. Žižek, *Against the double blackmail: Refugees, terror and other troubles with the neighbours*, Penguin UK, 2016.
- [789] K. Popper, *The Open Society and Its Enemies*. New One-Volume Edition, Princeton University Press, 2013.
- [790] B. Podobnik, I. S. Kirbis, M. Koprcina, H. Stanley, Emergence of the unified right- and left-wing populism—When radical societal changes become more important than ideology, *Physica A* 517 (2019) 459–474.
- [791] S. N. Durlauf, How can statistical mechanics contribute to social science?, *Proc. Natl. Acad. Sci. USA* 96 (1999) 10582–10584.
- [792] W. Weidlich, Physics and social science—the approach of synergetics, *Phys. Rep.* 204 (1991) 1–163.
- [793] W. Weidlich, *Sociodynamics: A systematic approach to mathematical modelling in the social sciences*, Courier Corporation, 2006.
- [794] S. Galam, Sociophysics: A review of Galam models, *Int. J. Mod. Phys. C* 19 (2008) 409–440.

- [795] K. Bansak, J. Ferwerda, J. Hainmueller, A. Dillon, D. Hangartner, D. Lawrence, J. Weinstein, Improving refugee integration through data-driven algorithmic assignment, *Science* 359 (2018) 325–329.
- [796] P. A. Van Lange, J. Joireman, C. D. Parks, E. Van Dijk, The psychology of social dilemmas: A review, *Organ. Behav. Hum. Decis. Process.* 120 (2013) 125–141.
- [797] C. Hauert, G. Szabó, Game theory and physics, *Am. J. Phys.* 73 (2005) 405–414.
- [798] M. E. J. Newman, D. J. Watts, Renormalization group analysis of the small-world network model, *Phys. Lett. A* 263 (1999) 341–346.
- [799] H. Hinrichsen, Non-equilibrium critical phenomena and phase transitions into absorbing states, *Adv. Phys.* 49 (2000) 815–958.
- [800] G. Szabó, C. Hauert, Phase transitions and volunteering in spatial public goods games, *Phys. Rev. Lett.* 89 (2002) 118101.
- [801] M. Perc, Phase transitions in models of human cooperation, *Phys. Lett. A* 380 (2016) 2803–2808.
- [802] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. Van Nes, M. Rietkerk, G. Sugihara, Early-warning signals for critical transitions, *Nature* 461 (2009) 53–59.
- [803] M. Scheffer, S. R. Carpenter, Catastrophic regime shifts in ecosystems: linking theory to observation, *Trends Ecol. Evol.* 18 (2003) 648–656.
- [804] M. Levy, Social phase transitions, *J. Econ. Behav. Organ.* 57 (2005) 71–87.
- [805] M. Levy, Stock market crashes as social phase transitions, *J. Econ. Dyn. Control* 32 (2008) 137–155.
- [806] Standard Eurobarometer 85—Spring 2016, 2016. Survey conducted by TNS opinion & social at the request of the European Commission, Directorate-General for Communication. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.

- [807] R. M. White, T. H. Geballe, Long range order in solids, Academic Press, 1979.
- [808] G. Bertotti, Hysteresis in magnetism: for physicists, materials scientists, and engineers, Academic Press, 1998.
- [809] B. Podobnik, D. Korošak, M. S. Klemen, A. Stožer, J. Dolenšek, M. S. Rupnik, P. C. Ivanov, P. Holme, M. Jusup, β -cells operate collectively to help maintain glucose homeostasis, *Biophys. J.* 118 (2020) 2588–2595.
- [810] U. Witt, Self-organization and economics—what is new?, *Struct. Chang. Econ. Dyn.* 8 (1997) 489–507.
- [811] R. Pfeifer, M. Lungarella, F. Iida, Self-organization, embodiment, and biologically inspired robotics, *Science* 318 (2007) 1088–1093.
- [812] D. Helbing, Social self-organization: Agent-based simulations and experiments to study emergent social behavior, Springer, 2012.
- [813] J. Park, Europe’s migrant crisis, 2015. Council on Foreign Relations. Available at: <https://www.cfr.org/backgrounder/europes-migration-crisis>. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [814] V. Mueller, C. Gray, K. Kosec, Heat stress increases long-term human migration in rural Pakistan, *Nat. Clim. Change* 4 (2014) 182–185.
- [815] R. Forst, Toleration, justice and reason, in: D. Castiglione, C. Mckinnon (Eds.), *The culture of toleration in diverse societies*, Manchester University Press, 2003, pp. 71–85.
- [816] J. L. Sullivan, J. Piereson, G. E. Marcus, *Political tolerance and American democracy*, University of Chicago Press, 1993.
- [817] R. M. Kunovich, Labor market competition and anti-immigrant sentiment: Occupations as contexts, *Int. Migr. Rev.* 47 (2013) 643–685.
- [818] T. Paas, V. Halapuu, Attitudes towards immigrants and the integration of ethnically diverse societies., *East. J. Eur. Stud.* 3 (2012) 161–176.

- [819] N. Esipova, J. Ray, Syrian refugees not welcome in Eastern Europe, 2017. Gallup World Poll 2016. Available at: <https://news.gallup.com/poll/209828/syrian-refugees-not-welcome-eastern-europe.aspx>. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [820] M. A. Nowak, K. Sigmund, Evolution of indirect reciprocity by image scoring, *Nature* 393 (1998) 573–577.
- [821] R. L. Riolo, M. D. Cohen, R. Axelrod, Evolution of cooperation without reciprocity, *Nature* 414 (2001) 441–443.
- [822] M. Semyonov, R. Raijman, A. Gorodzeisky, The rise of anti-foreigner sentiment in European societies, 1988-2000, *Am. Sociol. Rev.* 71 (2006) 426–449.
- [823] A. Gorodzeisky, M. Semyonov, Not only competitive threat but also racial prejudice: Sources of anti-immigrant attitudes in European societies, *Int. J. Public Opin. Res.* 28 (2016) 331–354.
- [824] M. Hooghe, R. Dassonneville, Explaining the Trump vote: The effect of racist resentment and anti-immigrant sentiments, *PS Political Sci. Politics* 51 (2018) 528–534.
- [825] B. Podobnik, M. Jusup, Z. Wang, H. E. Stanley, How fear of future outcomes affects social dynamics, *J. Stat. Phys.* 167 (2017) 1007–1019.
- [826] B. R. Chiswick, The effect of americanization on the earnings of foreign-born men, *J. Political Econ.* 86 (1978) 897–921.
- [827] G. J. Borjas, The economics of immigration, *J. Econ. Lit.* 32 (1994) 1667–1717.
- [828] F. Büchel, J. R. Frick, Immigrants’ economic performance across Europe—does immigration policy matter?, *Popul. Res. Policy Rev.* 24 (2005) 175–212.
- [829] R. Axelrod, The dissemination of culture: A model with local convergence and global polarization, *J. Confl. Resolut.* 41 (1997) 203–226.
- [830] J.-P. Onnela, F. Reed-Tsochas, Spontaneous emergence of social influence in online systems, *Proc. Natl. Acad. Sci. USA* 107 (2010) 18375–18380.

- [831] C. Castellano, M. Marsili, A. Vespignani, Nonequilibrium phase transition in a model for social influence, *Phys. Rev. Lett.* 85 (2000) 3536.
- [832] G. Weisbuch, G. Deffuant, F. Amblard, J.-P. Nadal, Meet, discuss, and segregate!, *Complexity* 7 (2002) 55–63.
- [833] K. Klemm, V. M. Eguíluz, R. Toral, M. San Miguel, Nonequilibrium transitions in complex networks: A model of social interaction, *Phys. Rev. E* 67 (2003) 026120.
- [834] K. Klemm, V. M. Eguíluz, R. Toral, M. San Miguel, Global culture: A noise-induced transition in finite systems, *Phys. Rev. E* 67 (2003) 045101.
- [835] G. Weisbuch, Bounded confidence and social networks, *Eur. Phys. J. B* 38 (2004) 339–343.
- [836] E. Lee, P. Holme, S. H. Lee, Modeling the dynamics of dissent, *Physica A* 486 (2017) 262–272.
- [837] J. S. Juul, M. A. Porter, Hipsters on networks: How a minority group of individuals can lead to an antiestablishment majority, *Phys. Rev. E* 99 (2019) 022313.
- [838] J. D. Touboul, The hipster effect: When anti-conformists all look the same, *Discrete Contin. Dyn. Syst. Ser. B* 24 (2019) 4379.
- [839] P. Dittrich, F. Liljeros, A. Soulier, W. Banzhaf, Spontaneous group formation in the seceder model, *Phys. Rev. Lett.* 84 (2000) 3205.
- [840] P. Dittrich, W. Banzhaf, Survival of the unfittest?—the seceder model and its fitness landscape, in: J. Kelemen, P. Sosík (Eds.), *Advances in Artificial Life*, Springer, 2001, pp. 100–109.
- [841] A. Grönlund, P. Holme, Networking the seceder model: Group formation in social and economic systems, *Phys. Rev. E* 70 (2004) 036108.
- [842] P. Holme, A. Grönlund, Modeling the dynamics of youth subcultures, *J. Artif. Soc. Soc. Simul.* 8 (2005) 3.
- [843] C. Mudde, The populist zeitgeist, *Gov. Oppos.* 39 (2004) 541–563.

- [844] H.-G. Betz, The two faces of radical right-wing populism in Western Europe, *Rev. Politics* 55 (1993) 663–685.
- [845] D. Oesch, Explaining workers’ support for right-wing populist parties in Western Europe: Evidence from Austria, Belgium, France, Norway, and Switzerland, *Int. Political Sci. Rev.* 29 (2008) 349–373.
- [846] K. Arzheimer, E. Carter, Political opportunity structures and right-wing extremist party success, *Eur. J. Political Res.* 45 (2006) 419–443.
- [847] T. Bjørklund, Unemployment and the radical right in Scandinavia: beneficial or non-beneficial for electoral support?, *Comp. Eur. Politics* 5 (2007) 245–263.
- [848] J. M. Smith, Does crime pay? Issue ownership, political opportunity, and the populist right in Western Europe, *Comp. Political Stud.* 43 (2010) 1471–1498.
- [849] C. Lawton, R. Ackrill, Hard evidence: how areas with low immigration voted mainly for Brexit, 2016. Available at: <https://theconversation.com/hard-evidence-how-areas-with-low-immigration-voted-mainly-for-brexit-62138>. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [850] G. Deffuant, D. Neau, F. Amblard, G. Weisbuch, Mixing beliefs among interacting agents, *Adv. Complex Syst.* 3 (2000) 87–98.
- [851] A. I. Abramowitz, K. L. Saunders, Is polarization a myth?, *J. Politics* 70 (2008) 542–555.
- [852] S. Saavedra, J. Efstathiou, F. Reed-Tsochas, Identifying the underlying structure and dynamic interactions in a voting network, *Physica A* 377 (2007) 672–688.
- [853] L. Jasny, J. Waggle, D. R. Fisher, An empirical examination of echo chambers in us climate policy networks, *Nat. Clim. Change* 5 (2015) 782–786.
- [854] H. T. Williams, J. R. McMurray, T. Kurz, F. H. Lambert, Network analysis reveals open forums and echo chambers in social media discussions of climate change, *Glob. Environ. Change* 32 (2015) 126–138.

- [855] A. Boutyline, R. Willer, The social structure of political echo chambers: Variation in ideological homophily in online networks, *Political Psychol.* 38 (2017) 551–569.
- [856] N. Usher, J. Holcomb, J. Littman, Twitter makes it worse: Political journalists, gendered echo chambers, and the amplification of gender bias, *Int. J. Press/Politics* 23 (2018) 324–344.
- [857] M. Cinelli, G. D. F. Morales, A. Galeazzi, W. Quattrociocchi, M. Starnini, The echo chamber effect on social media, *Proc. Natl. Acad. Sci. USA* 118 (2021) e2023301118.
- [858] S. Flaxman, S. Goel, J. M. Rao, Filter bubbles, echo chambers, and online news consumption, *Public Opin. Q.* 80 (2016) 298–320.
- [859] E. Ben-Naim, P. L. Krapivsky, S. Redner, Bifurcations and patterns in compromise processes, *Physica D* 183 (2003) 190–204.
- [860] F. Amblard, G. Deffuant, The role of network topology on extremism propagation with the relative agreement opinion dynamics, *Physica A* 343 (2004) 725–738.
- [861] D. Stauffer, H. Meyer-Ortmanns, Simulation of consensus model of Deffuant et al. on a Barabasi-Albert network, *Int. J. Mod. Phys. C* 15 (2004) 241–246.
- [862] E. Ben-Naim, Opinion dynamics: rise and fall of political parties, *EPL (Europhys. Lett.)* 69 (2005) 671.
- [863] P. Holme, M. E. J. Newman, Nonequilibrium phase transition in the coevolution of networks and opinions, *Phys. Rev. E* 74 (2006) 056108.
- [864] T. Gross, B. Blasius, Adaptive coevolutionary networks: a review, *J. R. Soc. Interface* 5 (2008) 259–271.
- [865] B. Kozma, A. Barrat, Consensus formation on adaptive networks, *Phys. Rev. E* 77 (2008) 016102.
- [866] F. Baumann, P. Lorenz-Spreen, I. M. Sokolov, M. Starnini, Modeling echo chambers and polarization dynamics in social networks, *Phys. Rev. Lett.* 124 (2020) 048301.

- [867] K. Suchecki, V. M. Eguíluz, M. San Miguel, Voter model dynamics in complex networks: Role of dimensionality, disorder, and degree distribution, *Phys. Rev. E* 72 (2005) 036132.
- [868] M. Mobilia, A. Petersen, S. Redner, On the role of zealotry in the voter model, *J. Stat. Mech. Theory Exp.* 2007 (2007) P08029.
- [869] N. Masuda, N. Gibert, S. Redner, Heterogeneous voter models, *Phys. Rev. E* 82 (2010) 010103.
- [870] J. Fernández-Gracia, K. Suchecki, J. J. Ramasco, M. San Miguel, V. M. Eguíluz, Is the voter model a model for voters?, *Phys. Rev. Lett.* 112 (2014) 158701.
- [871] S. Redner, Reality-inspired voter models: A mini-review, *C. R. Phys.* 20 (2019) 275–292.
- [872] L. Steels, A self-organizing spatial vocabulary, *Artif. Life* 2 (1995) 319–332.
- [873] A. Baronchelli, M. Felici, V. Loreto, E. Caglioti, L. Steels, Sharp transition towards shared vocabularies in multi-agent systems, *J. Stat. Mech. Theory Exp.* 2006 (2006) P06014.
- [874] V. Loreto, L. Steels, Emergence of language, *Nat. Phys.* 3 (2007) 758–760.
- [875] A. Baronchelli, V. Loreto, L. Steels, In-depth analysis of the Naming Game dynamics: the homogeneous mixing case, *Int. J. Mod. Phys. C* 19 (2008) 785–812.
- [876] Q. Lu, G. Korniss, B. K. Szymanski, The naming game in social networks: community formation and consensus engineering, *J. Econ. Interact. Coord.* 4 (2009) 221–235.
- [877] E. DeGiuli, Random language model, *Phys. Rev. Lett.* 122 (2019) 128301.
- [878] A. Vespignani, Modelling dynamical processes in complex socio-technical systems, *Nat. Phys.* 8 (2012) 32–39.

- [879] G. Neumann, T. Noda, Y. Kawaoka, Emergence and pandemic potential of swine-origin h1n1 influenza virus, *Nature* 459 (2009) 931–939.
- [880] D. Liu, W. Shi, Y. Shi, D. Wang, H. Xiao, W. Li, Y. Bi, Y. Wu, X. Li, J. Yan, et al., Origin and diversity of novel avian influenza A H7N9 viruses causing human infection: phylogenetic, structural, and coalescent analyses, *Lancet* 381 (2013) 1926–1932.
- [881] A. J. Kucharski, W. J. Edmunds, Case fatality rate for Ebola virus disease in west Africa, *Lancet* 384 (2014) 1260.
- [882] D. Corti, J. Zhao, M. Pedotti, L. Simonelli, S. Agnihothram, C. Fett, B. Fernandez-Rodriguez, M. Foglierini, G. Agatic, F. Vanzetta, et al., Prophylactic and postexposure efficacy of a potent human monoclonal antibody against MERS coronavirus, *Proc. Natl. Acad. Sci. USA* 112 (2015) 10473–10478.
- [883] W. O. Kermack, A. G. McKendrick, A contribution to the mathematical theory of epidemics, *Proc. R. Soc. Lond. A* 115 (1927) 700–721.
- [884] H. Abbey, An examination of the Reed-Frost theory of epidemics, *Hum. Biol.* 24 (1952) 201–233.
- [885] M. Y. Li, Z. Shuai, Global stability of an epidemic model in a patchy environment, *Can. Appl. Math. Q.* 17 (2009) 175–187.
- [886] R. Guimera, S. Mossa, A. Turtschi, L. N. Amaral, The worldwide air transportation network: Anomalous centrality, community structure, and cities’ global roles, *Proc. Natl. Acad. Sci. USA* 102 (2005) 7794–7799.
- [887] L. Hufnagel, D. Brockmann, T. Geisel, Forecast and control of epidemics in a globalized world, *Proc. Natl. Acad. Sci. USA* 101 (2004) 15124–15129.
- [888] V. Colizza, A. Barrat, M. Barthélemy, A. Vespignani, The role of the airline transportation network in the prediction and predictability of global epidemics, *Proc. Natl. Acad. Sci. USA* 103 (2006) 2015–2020.
- [889] P. Bajardi, C. Poletto, J. J. Ramasco, M. Tizzoni, V. Colizza, A. Vespignani, Human mobility networks, travel restrictions, and the global spread of 2009 H1N1 pandemic, *PLOS ONE* 6 (2011) e16591.

- [890] L. Wang, J. T. Wu, Characterizing the dynamics underlying global spread of epidemics, *Nat. Commun.* 9 (2018) 218.
- [891] D. Balcan, V. Colizza, B. Gonçalves, H. Hu, J. J. Ramasco, A. Vespignani, Multiscale mobility networks and the spatial spreading of infectious diseases, *Proc. Natl. Acad. Sci. USA* 106 (2009) 21484–21489.
- [892] C. Poletto, M. Tizzoni, V. Colizza, Heterogeneous length of stay of hosts’ movements and spatial epidemic spread, *Sci. Rep.* 2 (2012) 476.
- [893] M. Chinazzi, J. T. Davis, M. Ajelli, C. Gioannini, M. Litvinova, S. Merler, A. P. y Piontti, K. Mu, L. Rossi, K. Sun, et al., The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak, *Science* 368 (2020) 395–400.
- [894] D. Soriano-Paños, L. Lotero, A. Arenas, J. Gómez-Gardeñes, Spreading processes in multiplex metapopulations containing different mobility networks, *Phys. Rev. X* 8 (2018) 031039.
- [895] T. Bedford, S. Riley, I. G. Barr, S. Broor, M. Chadha, N. J. Cox, R. S. Daniels, C. P. Gunasekaran, A. C. Hurt, A. Kelso, et al., Global circulation patterns of seasonal influenza viruses vary with antigenic drift, *Nature* 523 (2015) 217–220.
- [896] M. C. Gonzalez, C. A. Hidalgo, A.-L. Barabasi, Understanding individual human mobility patterns, *nature* 453 (2008) 779–782.
- [897] A. Wesolowski, N. Eagle, A. J. Tatem, D. L. Smith, A. M. Noor, R. W. Snow, C. O. Buckee, Quantifying the impact of human mobility on malaria, *Science* 338 (2012) 267–270.
- [898] S. Lai, N. W. Ruktanonchai, L. Zhou, O. Prosper, W. Luo, J. R. Floyd, A. Wesolowski, M. Santillana, C. Zhang, X. Du, et al., Effect of non-pharmaceutical interventions to contain COVID-19 in China, *Nature* 585 (2020) 410–413.
- [899] L. A. Castro, N. Generous, W. Luo, A. Pastore y Piontti, K. Martinez, M. F. Gomes, D. Osthus, G. Fairchild, A. Ziemann, A. Vespignani, et al., Using heterogeneous data to identify signatures of dengue outbreaks at fine spatio-temporal scales across Brazil, *PLOS Negl. Trop. Dis.* 15 (2021) e0009392.

- [900] P. Nouvellet, S. Bhatia, A. Cori, K. E. Ainslie, M. Baguelin, S. Bhatt, A. Boonyasiri, N. F. Brazeau, L. Cattarino, L. V. Cooper, et al., Reduction in mobility and COVID-19 transmission, *Nat. Commun.* 12 (2021) 1090.
- [901] X.-K. Xu, L. Wang, S. Pei, Multiscale mobility explains differential associations between the gross domestic product and COVID-19 transmission in Chinese cities, *J. Travel Med.* 28 (2021) taaa236.
- [902] M. Salathé, Digital epidemiology: what is it, and where is it going?, *Life Sci. Soc. Policy* 14 (2018) 1–5.
- [903] G. Cervellin, I. Comelli, G. Lippi, Is Google Trends a reliable tool for digital epidemiology? Insights from different clinical settings, *J. Epidemiol. Glob. Health* 7 (2017) 185–189.
- [904] L. Bengtsson, X. Lu, A. Thorson, R. Garfield, J. Von Schreeb, Improved response to disasters and outbreaks by tracking population movements with mobile phone network data: a post-earthquake geospatial study in Haiti, *PLOS Med.* 8 (2011) e1001083.
- [905] P. M. Polgreen, Y. Chen, D. M. Pennock, F. D. Nelson, R. A. Weinstein, Using internet searches for influenza surveillance, *Clin. Infect. Dis.* 47 (2008) 1443–1448.
- [906] D. Lazer, R. Kennedy, G. King, A. Vespignani, The parable of Google Flu: traps in big data analysis, *Science* 343 (2014) 1203–1205.
- [907] M. Santillana, D. W. Zhang, B. M. Althouse, J. W. Ayers, What can digital disease detection learn from (an external revision to) Google Flu Trends?, *Am. J. Prev. Med.* 47 (2014) 341–347.
- [908] S. Yang, M. Santillana, S. C. Kou, Accurate estimation of influenza epidemics using Google search data via ARGO, *Proc. Natl. Acad. Sci. USA* 112 (2015) 14473–14478.
- [909] F. S. Lu, M. W. Hattab, C. L. Clemente, M. Biggerstaff, M. Santillana, Improved state-level influenza nowcasting in the United States leveraging Internet-based data and network approaches, *Nat. Commun.* 10 (2019) 147.

- [910] L. Clemente, F. Lu, M. Santillana, Improved real-time influenza surveillance: using internet search data in eight Latin American countries, *JMIR Public Health Surveill.* 5 (2019) e12214.
- [911] B. M. Althouse, Y. Y. Ng, D. A. T. Cummings, Prediction of dengue incidence using search query surveillance, *PLOS Negl. Trop. Dis.* 5 (2011) e1258.
- [912] E. H. Chan, V. Sahai, C. Conrad, J. S. Brownstein, Using web search query data to monitor dengue epidemics: a new model for neglected tropical disease surveillance, *PLOS Negl. Trop. Dis.* 5 (2011) e1206.
- [913] S. Alasaad, War diseases revealed by the social media: massive leishmaniasis outbreak in the Syrian Spring, *Parasites Vectors* 6 (2013) 94.
- [914] A. J. Ocampo, R. Chunara, J. S. Brownstein, Using search queries for malaria surveillance, Thailand, *Malar. J.* 12 (2013) 390.
- [915] G. J. Milinovich, S. M. Avril, A. C. Clements, J. S. Brownstein, S. Tong, W. Hu, Using internet search queries for infectious disease surveillance: screening diseases for suitability, *BMC Infect. Dis.* 14 (2014) 690.
- [916] J. P. Messina, O. J. Brady, D. M. Pigott, J. S. Brownstein, A. G. Hoen, S. I. Hay, A global compendium of human dengue virus occurrence, *Sci. Data* 1 (2014) 140004.
- [917] R. T. Gluskin, M. A. Johansson, M. Santillana, J. S. Brownstein, Evaluation of Internet-based dengue query data: Google Dengue Trends, *PLOS Negl. Trop. Dis.* 8 (2014) e2713.
- [918] S. Yang, S. C. Kou, F. Lu, J. S. Brownstein, N. Brooke, M. Santillana, Advances in using Internet searches to track dengue, *PLOS Comput. Biol.* 13 (2017) e1005607.
- [919] R. A. Strauss, J. S. Castro, R. Reintjes, J. R. Torres, Google dengue trends: An indicator of epidemic behavior. The Venezuelan Case, *Int. J. Med. Inform.* 104 (2017) 26–30.

- [920] H. T. Ho, T. M. Carvajal, J. R. Bautista, J. D. R. Capistrano, K. M. Viacrusis, L. F. T. Hernandez, K. Watanabe, Using Google trends to examine the spatio-temporal incidence and behavioral patterns of dengue disease: a case study in metropolitan Manila, Philippines, *Trop. Med. Infect. Dis.* 3 (2018) 118.
- [921] F. Gesualdo, G. Stilo, M. V. Gonfiantini, E. Pandolfi, P. Velardi, A. E. Tozzi, Influenza-like illness surveillance on Twitter through automated learning of naïve language, *PLOS ONE* 8 (2013) e82489.
- [922] D. A. Broniatowski, M. J. Paul, M. Dredze, National and local influenza surveillance through Twitter: an analysis of the 2012-2013 influenza epidemic, *PLOS ONE* 8 (2013) e83672.
- [923] C. de Almeida Marques-Toledo, C. M. Degener, L. Vinhal, G. Coelho, W. Meira, C. T. Codeço, M. M. Teixeira, Dengue prediction by the web: Tweets are a useful tool for estimating and forecasting Dengue at country and city level, *PLOS Negl. Trop. Dis.* 11 (2017) e0005729.
- [924] M. S. Deiner, C. Fathy, J. Kim, K. Niemeyer, D. Ramirez, S. F. Ackley, F. Liu, T. M. Lietman, T. C. Porco, Facebook and Twitter vaccine sentiment in response to measles outbreaks, *Health Inform. J.* 25 (2019) 1116–1132.
- [925] T. Hu, S. Wang, W. Luo, Y. Yan, M. Zhang, X. Huang, R. Liu, K. Ly, V. Kacker, Z. Li, Revealing public opinion towards COVID-19 vaccines using Twitter data in the United States: a spatiotemporal perspective, *J. Med. Internet Res.* 23 (2021) e30854.
- [926] M. Salathé, S. Khandelwal, Assessing vaccination sentiments with online social media: implications for infectious disease dynamics and control, *PLOS Comput. Biol.* 7 (2011) e1002199.
- [927] M. Salathé, D. Q. Vu, S. Khandelwal, D. R. Hunter, The dynamics of health behavior sentiments on a large online social network, *EPJ Data Sci.* 2 (2013) 4.
- [928] J. Bian, U. Topaloglu, F. Yu, Towards large-scale twitter mining for drug-related adverse events, in: C. C. Yang, H. Chen, H. Wactlar, C. Combi, X. Tang (Eds.), *Proceedings of the 2012 international*

workshop on Smart health and wellbeing, Association for Computing Machinery, 2012, pp. 25–32.

- [929] R. Jurdak, K. Zhao, J. Liu, M. AbouJaoude, M. Cameron, D. Newth, Understanding human mobility from Twitter, *PLOS ONE* 10 (2015) e0131469.
- [930] J. I. Blanford, Z. Huang, A. Savelyev, A. M. MacEachren, Geo-located tweets. enhancing mobility maps and capturing cross-border movement, *PLOS ONE* 10 (2015) e0129202.
- [931] Facebook Data for Good Annual Report 2020, 2020. Available at: <https://dataforgood.fb.com/wp-content/uploads/2021/01/Facebook-Data-for-Good-2020-Annual-Report-1.pdf>. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [932] T. Kuchler, D. Russel, J. Stroebel, JUE Insight: The geographic spread of COVID-19 correlates with the structure of social networks as measured by Facebook, *J. Urban Econ.* (2021) 103314.
- [933] A. Spelta, A. Flori, F. Pierri, G. Bonaccorsi, F. Pammolli, After the lockdown: simulating mobility, public health and economic recovery scenarios, *Sci. Rep.* 10 (2020) 16950.
- [934] P. Maas, S. Iyer, A. Gros, W. Park, L. McGorman, C. Nayak, P. A. Dow, Facebook Disaster Maps: Aggregate insights for crisis response & recovery., in: Z. Franco, J. J. González, J. H. Canós (Eds.), *Proceedings of the 16th International Conference on Information Systems for Crisis Response And Management*, International Association for Information Systems for Crisis Response and Management, 2019, pp. 836–847.
- [935] S. Cook, C. Conrad, A. L. Fowlkes, M. H. Mohebbi, Assessing Google flu trends performance in the United States during the 2009 influenza virus A(H1N1) pandemic, *PLOS ONE* 6 (2011) e23610.
- [936] D. Butler, When Google got flu wrong: US outbreak foxes a leading web-based method for tracking seasonal flu, *Nature* 494 (2013) 155–157.

- [937] D. R. Olson, K. J. Konty, M. Paladini, C. Viboud, L. Simonsen, Re-assessing Google Flu Trends data for detection of seasonal and pandemic influenza: a comparative epidemiological study at three geographic scales, *PLOS Comput. Biol.* 9 (2013) e1003256.
- [938] M. Santillana, A. T. Nguyen, M. Dredze, M. J. Paul, E. O. Nsoesie, J. S. Brownstein, Combining search, social media, and traditional data sources to improve influenza surveillance, *PLOS Comput. Biol.* 11 (2015) e1004513.
- [939] M. S. Smolinski, A. W. Crawley, K. Baltrusaitis, R. Chunara, J. M. Olsen, O. Wójcik, M. Santillana, A. Nguyen, J. S. Brownstein, Flu Near You: Crowdsourced symptom reporting spanning 2 influenza seasons, *Am. J. Public Health* 105 (2015) 2124–2130.
- [940] M. Santillana, Editorial commentary: Perspectives on the future of internet search engines and biosurveillance systems, *Clin. Infect. Dis.* 64 (2016) 42–43.
- [941] S. F. McGough, J. S. Brownstein, J. B. Hawkins, M. Santillana, Forecasting Zika incidence in the 2016 Latin America outbreak combining traditional disease surveillance with search, social media, and news report data, *PLOS Negl. Trop. Dis.* 11 (2017) e0005295.
- [942] J. S. Brownstein, C. C. Freifeld, B. Y. Reis, K. D. Mandl, Surveillance Sans Frontieres: Internet-based emerging infectious disease intelligence and the HealthMap project, *PLOS Med.* 5 (2008) e151.
- [943] K. Baltrusaitis, J. S. Brownstein, S. V. Scarpino, E. Bakota, A. W. Crawley, G. Conidi, J. Gunn, J. Gray, A. Zink, M. Santillana, Comparison of crowd-sourced, electronic health records based, and traditional health-care based influenza-tracking systems at multiple spatial resolutions in the United States of America, *BMC Infect. Dis.* 18 (2018) 403.
- [944] D. Liu, L. Clemente, C. Poirier, X. Ding, M. Chinazzi, J. Davis, A. Vespignani, M. Santillana, Real-time forecasting of the COVID-19 outbreak in Chinese provinces: machine learning approach using novel digital data and estimates from mechanistic models, *J. Med. Internet Res.* 22 (2020) e20285.

- [945] F. S. Lu, S. Hou, K. Baltrusaitis, M. Shah, J. Leskovec, J. Hawkins, J. Brownstein, G. Conidi, J. Gunn, J. Gray, et al., Accurate influenza monitoring and forecasting using novel internet data streams: a case study in the Boston Metropolis, *JMIR Public Health Surveill.* 4 (2018) e4.
- [946] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, A. Vespignani, Epidemic processes in complex networks, *Rev. Mod. Phys.* 87 (2015) 925.
- [947] R. Pastor-Satorras, A. Vespignani, Epidemic spreading in scale-free networks, *Phys. Rev. Lett.* 86 (2001) 3200.
- [948] C. Castellano, R. Pastor-Satorras, Thresholds for epidemic spreading in networks, *Phys. Rev. Lett.* 105 (2010) 218701.
- [949] R. Cohen, S. Havlin, D. Ben-Avraham, Efficient immunization strategies for computer networks and populations, *Phys. Rev. Lett.* 91 (2003) 247901.
- [950] V. Colizza, A. Vespignani, Invasion threshold in heterogeneous metapopulation networks, *Phys. Rev. Lett.* 99 (2007) 148701.
- [951] V. Colizza, A. Vespignani, Epidemic modeling in metapopulation systems with heterogeneous coupling pattern: Theory and simulations, *J. Theor. Biol.* 251 (2008) 450–467.
- [952] D. Balcan, A. Vespignani, Phase transitions in contagion processes mediated by recurrent mobility patterns, *Nat. Phys.* 7 (2011) 581–586.
- [953] D. Balcan, A. Vespignani, Invasion threshold in structured populations with recurrent mobility patterns, *J. Theor. Biol.* 293 (2012) 87–100.
- [954] C. Fraser, C. A. Donnelly, S. Cauchemez, W. P. Hanage, M. D. Van Kerkhove, T. D. Hollingsworth, J. Griffin, R. F. Baggaley, H. E. Jenkins, E. J. Lyons, et al., Pandemic potential of a strain of influenza A(H1N1): early findings, *Science* 324 (2009) 1557–1561.
- [955] Z. Du, L. Wang, S. Cauchemez, X. Xu, X. Wang, B. J. Cowling, L. A. Meyers, Risk for transportation of coronavirus disease from Wuhan to other cities in China, *Emerg. Infect. Dis.* 26 (2020) 1049–1052.

- [956] J. T. Wu, K. Leung, G. M. Leung, Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study, *Lancet* 395 (2020) 689–697.
- [957] A. Gautreau, A. Barrat, M. Barthelemy, Global disease spread: statistics and estimation of arrival times, *J. Theor. Biol.* 251 (2008) 509–522.
- [958] G. S. Tomba, J. Wallinga, A simple explanation for the low impact of border control as a countermeasure to the spread of an infectious disease, *Math. Biosci.* 214 (2008) 70–72.
- [959] D. Brockmann, D. Helbing, The hidden geometry of complex, network-driven contagion phenomena, *Science* 342 (2013) 1337–1342.
- [960] S. M. Ross, *Stochastic processes*, John Wiley & Sons, 1996.
- [961] N. D. Grubaugh, J. T. Ladner, M. U. Kraemer, G. Dudas, A. L. Tan, K. Gangavarapu, M. R. Wiley, S. White, J. Thézé, D. M. Magnani, et al., Genomic epidemiology reveals multiple introductions of Zika virus into the United States, *Nature* 546 (2017) 401–405.
- [962] A. da Silva Filipe, J. G. Shepherd, T. Williams, J. Hughes, E. Aranday-Cortes, P. Asamaphan, S. Ashraf, C. Balcazar, K. Brunker, A. Campbell, et al., Genomic epidemiology reveals multiple introductions of SARS-CoV-2 from mainland Europe into Scotland, *Nat. Microbiol.* 6 (2020) 112–122.
- [963] T. M. Cover, J. A. Thomas, *Elements of information theory*, 2nd edition, John Wiley & Sons, 2006.
- [964] R. van de Schoot, S. Depaoli, R. King, B. Kramer, K. Märten, M. G. Tadesse, M. Vannucci, A. Gelman, D. Veen, J. Willemsen, et al., Bayesian statistics and modelling, *Nat. Rev. Methods Primers* 1 (2021) 1.
- [965] B. Lambert, *A student’s guide to Bayesian statistics*, Sage, 2018.
- [966] A. Gelman, D. Simpson, M. Betancourt, The prior can often only be understood in the context of the likelihood, *Entropy* 19 (2017) 555.

- [967] N. P. Lemoine, Moving beyond noninformative priors: why and how to choose weakly informative priors in Bayesian analyses, *Oikos* 128 (2019) 912–928.
- [968] D. Balcan, H. Hu, B. Goncalves, P. Bajardi, C. Poletto, J. J. Ramasco, D. Paolotti, N. Perra, M. Tizzoni, W. Van den Broeck, et al., Seasonal transmission potential and activity peaks of the new influenza A(H1N1): a Monte Carlo likelihood analysis based on human mobility, *BMC Med.* 7 (2009) 45.
- [969] Q. Li, X. Guan, P. Wu, X. Wang, L. Zhou, Y. Tong, R. Ren, K. S. Leung, E. H. Lau, J. Y. Wong, et al., Early transmission dynamics in Wuhan, China, of novel coronavirus—infected pneumonia, *N. Engl. J. Med.* 382 (2020) 1199–1207.
- [970] N. Imai, I. Dorigatti, A. Cori, C. Donnelly, S. Riley, N. M. Ferguson, Estimating the potential total number of novel Coronavirus cases in Wuhan City, China, 2020. Report by Imperial College London. Available at: <https://spiral.imperial.ac.uk/handle/10044/1/77150>. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [971] R. Li, S. Pei, B. Chen, Y. Song, T. Zhang, W. Yang, J. Shaman, Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV-2), *Science* 368 (2020) 489–493.
- [972] S. Sanche, Y. T. Lin, C. Xu, E. Romero-Severson, N. Hengartner, R. Ke, High contagiousness and rapid spread of severe acute respiratory syndrome coronavirus 2, *Emerg. Infect. Dis.* 26 (2020) 1470–1477.
- [973] W. Luo, P. Gao, S. Cassels, A large-scale location-based social network to understanding the impact of human geo-social interaction patterns on vaccination strategies in an urbanized area, *Comput. Environ. Urban Syst.* 72 (2018) 78–87.
- [974] L. Yin, H. Zhang, Y. Li, K. Liu, T. Chen, W. Luo, S. Lai, Y. Li, X. Tang, L. Ning, et al., A data driven agent-based model that recommends non-pharmaceutical interventions to suppress Coronavirus disease 2019 resurgence in megacities, *J. R. Soc. Interface* 18 (2021) 20210112.

- [975] S. Chang, E. Pierson, P. W. Koh, J. Gerardin, B. Redbird, D. Grusky, J. Leskovec, Mobility network models of COVID-19 explain inequities and inform reopening, *Nature* 589 (2021) 82–87.
- [976] J. Mossong, N. Hens, M. Jit, P. Beutels, K. Auranen, R. Mikolajczyk, M. Massari, S. Salmaso, G. S. Tomba, J. Wallinga, et al., Social contacts and mixing patterns relevant to the spread of infectious diseases, *PLOS Med.* 5 (2008) e74.
- [977] K. Prem, A. R. Cook, M. Jit, Projecting social contact matrices in 152 countries using contact surveys and demographic data, *PLOS Comput. Biol.* 13 (2017) e1005697.
- [978] D. Mistry, M. Litvinova, A. P. y Piontti, M. Chinazzi, L. Fumanelli, M. F. Gomes, S. A. Haque, Q.-H. Liu, K. Mu, X. Xiong, et al., Inferring high-resolution human mixing patterns for disease modeling, *Nat. Commun.* 12 (2021) 323.
- [979] M. O’Driscoll, G. R. Dos Santos, L. Wang, D. A. Cummings, A. S. Azman, J. Paireau, A. Fontanet, S. Cauchemez, H. Salje, Age-specific mortality and immunity patterns of SARS-CoV-2, *Nature* 590 (2021) 140–145.
- [980] D. C. Adam, P. Wu, J. Y. Wong, E. H. Lau, T. K. Tsang, S. Cauchemez, G. M. Leung, B. J. Cowling, Clustering and super-spreading potential of SARS-CoV-2 infections in Hong Kong, *Nat. Med.* 26 (2020) 1714–1719.
- [981] F. Wong, J. J. Collins, Evidence that coronavirus superspreading is fat-tailed, *Proc. Natl. Acad. Sci. USA* 117 (2020) 29416–29418.
- [982] J. O. Lloyd-Smith, S. J. Schreiber, P. E. Kopp, W. M. Getz, Super-spreading and the effect of individual variation on disease emergence, *Nature* 438 (2005) 355–359.
- [983] N. Perra, Non-pharmaceutical interventions during the COVID-19 pandemic: A review, *Phys. Rep.* 913 (2021) 1–52.
- [984] J. Zhang, M. Litvinova, Y. Liang, Y. Wang, W. Wang, S. Zhao, Q. Wu, S. Merler, C. Viboud, A. Vespignani, et al., Changes in

contact patterns shape the dynamics of the COVID-19 outbreak in China, *Science* 368 (2020) 1481–1486.

- [985] S. T. Ali, L. Wang, E. H. Lau, X.-K. Xu, Z. Du, Y. Wu, G. M. Leung, B. J. Cowling, Serial interval of SARS-CoV-2 was shortened over time by nonpharmaceutical interventions, *Science* 369 (2020) 1106–1109.
- [986] D. Paolotti, A. Carnahan, V. Colizza, K. Eames, J. Edmunds, G. Gomes, C. Koppeschaar, M. Rehn, R. Smallenburg, C. Turbelin, et al., Web-based participatory surveillance of infectious diseases: the Influenzanet participatory surveillance experience, *Clin. Microbiol. Infect.* 20 (2014) 17–21.
- [987] J. J. Thomas, K. A. Cook, A visual analytics agenda, *IEEE Comput. Graph. Appl.* 26 (2006) 10–13.
- [988] W. Luo, Visual analytics of geo-social interaction patterns for epidemic control, *Int. J. Health Geogr.* 15 (2016) 28.
- [989] C. N. Waters, J. Zalasiewicz, C. Summerhayes, A. D. Barnosky, C. Poirier, A. Gałuszka, A. Cearreta, M. Edgeworth, E. C. Ellis, M. Ellis, et al., The Anthropocene is functionally and stratigraphically distinct from the Holocene, *Science* 351 (2016) aad2622.
- [990] D. N. Pellow, Framing emerging environmental movement tactics: mobilizing consensus, demobilizing conflict, *Sociol. Forum* 14 (1999) 659–683.
- [991] M. J. Molina, F. S. Rowland, Stratospheric sink for chlorofluoromethanes: chlorine atom-catalysed destruction of ozone, *Nature* 249 (1974) 810–812.
- [992] J. C. Farman, B. G. Gardiner, J. D. Shanklin, Large losses of total ozone in Antarctica reveal seasonal ClO_x/NO_x interaction, *Nature* 315 (1985) 207–210.
- [993] Ozone Secretariat UNEP, Handbook for the Montreal protocol on substances that deplete the ozone layer, 7th edition, United Nations Environment Programme, 2006.

- [994] World Meteorological Organization (WMO), Scientific Assessment of Ozone Depletion: 2018, 2018. Global Ozone Research and Monitoring Project—Report No. 58. Available at: <https://ozone.unep.org/science/assessment/sap>. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [995] F. L. Hoffman, Deaths from lead poisoning, *Bulletin of the United States Bureau of Labor Statistics* 426 (1927) 1–45.
- [996] C. C. Patterson, Contaminated and natural lead environments of man, *Arch. Environ. Health* 11 (1965) 344–360.
- [997] J. Yoshinaga, Lead in the Japanese living environment, *Environ. Health Prev. Med.* 17 (2012) 433–443.
- [998] T. G. Deryabina, S. V. Kuchmel, L. L. Nagorskaya, T. G. Hinton, J. C. Beasley, A. Lerebours, J. T. Smith, Long-term census data reveal abundant wildlife populations at Chernobyl, *Curr. Biol.* 25 (2015) R824–R826.
- [999] J. De Bruijn, B. G. Hansen, S. Johansson, M. Luotamo, S. J. Munn, C. Musset, S. I. Olsen, H. Olsson, A. B. Paya-Perez, F. Pedersen, K. Rasmussen, B. Sokull-Kluttgen, Technical guidance document on risk assessment, European Communities, 2003. Available at: https://echa.europa.eu/documents/10162/16960216/tgdpart2_2ed_en.pdf. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [1000] S. Parasuraman, Toxicological screening, *J. Pharmacol. Pharmacother.* 2 (2011) 74–79.
- [1001] M. Witeska, B. Jezierska, The effects of environmental factors on metal toxicity to fish (review), *Fresenius Environ. Bull.* 12 (2003) 824–829.
- [1002] W. Wang, Factors affecting metal toxicity to (and accumulation by) aquatic organisms—Overview, *Environ. Int.* 13 (1987) 437–457.
- [1003] W. P. Norwood, U. Borgmann, D. G. Dixon, A. Wallace, Effects of metal mixtures on aquatic biota: a review of observations and methods, *Hum. Ecol. Risk Assess.* 9 (2003) 795–811.

- [1004] E. J. Llanos, W. Leal, D. H. Luu, J. Jost, P. F. Stadler, G. Restrepo, Exploration of the chemical space and its three historical regimes, *Proc. Natl. Acad. Sci. USA* 116 (2019) 12660–12665.
- [1005] T. S. Galloway, C. N. Lewis, Marine microplastics spell big problems for future generations, *Proc. Natl. Acad. Sci. USA* 113 (2016) 2331–2333.
- [1006] A. J. Vanbergen, the Insect Pollinators Initiative, Threats to an ecosystem service: pressures on pollinators, *Front. Ecol. Environ.* 11 (2013) 251–259.
- [1007] A. Barrick, A. Châtel, M. Bruneau, C. Mouneyrac, The role of high-throughput screening in ecotoxicology and engineered nanomaterials, *Environ. Toxicol. Chem.* 36 (2017) 1704–1714.
- [1008] C. A. Murphy, R. M. Nisbet, P. Antczak, N. Garcia-Reyero, A. Gergs, K. Lika, T. Mathews, E. B. Muller, D. Nacci, A. Peace, et al., Incorporating suborganismal processes into dynamic energy budget models for ecological risk assessment, *Integr. Environ. Assess. Manag.* 14 (2018) 615–624.
- [1009] G. T. Ankley, R. S. Bennett, R. J. Erickson, D. J. Hoff, M. W. Horning, R. D. Johnson, D. R. Mount, J. W. Nichols, C. L. Russom, P. K. Schmieder, et al., Adverse outcome pathways: a conceptual framework to support ecotoxicology research and risk assessment, *Environ. Toxicol. Chem.* 29 (2010) 730–741.
- [1010] G. M. Marques, S. Augustine, K. Lika, L. Pecquerie, T. Domingos, S. A. L. M. Kooijman, The AmP project: comparing species on the basis of dynamic energy budget parameters, *PLOS Comput. Biol.* 14 (2018) e1006100.
- [1011] L. Von Bertalanffy, Quantitative laws in metabolism and growth, *Q. Rev. Biol.* 32 (1957) 217–231.
- [1012] T. Jager, Some good reasons to ban EC_x and related concepts in ecotoxicology, *Environ. Sci. Technol.* 45 (2011) 8180–8181.

- [1013] T. Klanjscek, R. M. Nisbet, J. H. Priester, P. A. Holden, Modeling physiological processes that relate toxicant exposure and bacterial population dynamics, *PLOS ONE* 7 (2012) e26955.
- [1014] T. Klanjscek, R. M. Nisbet, J. H. Priester, P. A. Holden, Dynamic energy budget approach to modeling mechanisms of CdSe quantum dot toxicity, *Ecotoxicology* 22 (2013) 319–330.
- [1015] R. Fablet, L. Pecquerie, H. De Pontual, H. Høie, R. Millner, H. Mosegaard, S. A. L. M. Kooijman, Shedding light on fish otolith biomineralization using a bioenergetic approach, *PLOS ONE* 6 (2011) e27055.
- [1016] L. Pecquerie, R. Fablet, H. De Pontual, S. Bonhommeau, M. Alunno-Bruscia, P. Petitgas, S. A. L. M. Kooijman, Reconstructing individual food and growth histories from biogenic carbonates, *Mar. Ecol. Prog. Ser.* 447 (2012) 151–164.
- [1017] T. Klanjscek, E. B. Muller, P. A. Holden, R. M. Nisbet, Host-symbiont interaction model explains non-monotonic response of soybean growth and seed production to nano-CeO₂ exposure, *Environ. Sci. Technol.* 51 (2017) 4944–4950.
- [1018] T. Jager, C. Albert, T. G. Preuss, R. Ashauer, General unified threshold model of survival—a toxicokinetic-toxicodynamic framework for ecotoxicology, *Environ. Sci. Technol.* 45 (2011) 2529–2540.
- [1019] OECD, Current approaches in the statistical analysis of ecotoxicity data: A guidance to application, OECD Publishing, 2003. OECD Series on Testing and Assessment, No. 54. Available at: <https://doi.org/10.1787/9789264085275-en>.
- [1020] T. Jager, T. Vandenbrouck, J. Baas, W. M. De Coen, S. A. L. M. Kooijman, A biology-based approach for mixture toxicity of multiple endpoints over the life cycle, *Ecotoxicology* 19 (2010) 351–361.
- [1021] T. Jager, E. M. Gudmundsdóttir, N. Cedergreen, Dynamic modeling of sublethal mixture toxicity in the nematode *Caenorhabditis elegans*, *Environ. Sci. Technol.* 48 (2014) 7026–7033.

- [1022] A. Margerit, E. Gomez, R. Gilbin, Dynamic energy-based modeling of uranium and cadmium joint toxicity to *Caenorhabditis elegans*, *Chemosphere* 146 (2016) 405–412.
- [1023] T. Jager, B. T. Martin, E. I. Zimmer, DEBkiss or the quest for the simplest generic model of animal life history, *J. Theor. Biol.* 328 (2013) 9–18.
- [1024] G. Ledder, J. D. Logan, A. Joern, Dynamic energy budget models with size-dependent hazard rates, *J. Math. Biol.* 48 (2004) 605–622.
- [1025] B. T. Martin, E. I. Zimmer, V. Grimm, T. Jager, Dynamic Energy Budget theory meets individual-based modelling: a generic and accessible implementation, *Methods Ecol. Evol.* 3 (2012) 445–449.
- [1026] A. M. De Roos, Demographic analysis of continuous-time life-history models, *Ecol. Lett.* 11 (2008) 1–15.
- [1027] M. Beekman, M. Thompson, M. Jusup, Thermodynamic constraints and the evolution of parental provisioning in vertebrates, *Behav. Ecol.* 30 (2019) 583–591.
- [1028] T. Klanjscek, H. Caswell, M. G. Neubert, R. M. Nisbet, Integrating dynamic energy budgets into matrix population models, *Ecol. Model.* 196 (2006) 407–420.
- [1029] H. Ijima, M. Jusup, T. Takada, T. Akita, H. Matsuda, T. Klanjscek, Effects of environmental change and early-life stochasticity on Pacific bluefin tuna population growth, *Mar. Environ. Res.* 149 (2019) 18–26.
- [1030] A. M. De Roos, O. Diekmann, J. A. J. Metz, Studying the dynamics of structured population models: a versatile technique and its application to *Daphnia*, *Am. Nat.* 139 (1992) 123–147.
- [1031] A. M. De Roos, L. Persson, Physiologically structured models—from versatile technique to ecological theory, *Oikos* 94 (2001) 51–71.
- [1032] I. M. Smallegange, H. Caswell, M. E. M. Toorians, A. M. De Roos, Mechanistic description of population dynamics using dynamic energy budget theory incorporated into integral projection models, *Methods Ecol. Evol.* 8 (2017) 146–154.

- [1033] C. M. Hunter, H. Caswell, The use of the vec-permutation matrix in spatial matrix population models, *Ecol. Model.* 188 (2005) 15–21.
- [1034] S. J. Phillips, R. P. Anderson, R. E. Schapire, Maximum entropy modeling of species geographic distributions, *Ecol. Model.* 190 (2006) 231–259.
- [1035] S. J. Phillips, R. P. Anderson, M. Dudík, R. E. Schapire, M. E. Blair, Opening the black box: An open-source release of Maxent, *Ecography* 40 (2017) 887–893.
- [1036] S. Faurby, D. Silvestro, L. Werdelin, A. Antonelli, Brain expansion in early hominins predicts carnivore extinctions in East Africa, *Ecol. Lett.* 23 (2020) 537–544.
- [1037] C. M. Kennedy, J. R. Oakleaf, D. M. Theobald, S. Baruch-Mordo, J. Kiesecker, Managing the middle: A shift in conservation priorities based on the global human modification gradient, *Glob. Change Biol.* 25 (2019) 811–826.
- [1038] E. Belle, N. Kingston, N. Burgess, T. Sandwith, N. Ali, K. MacKinnon, Protected Planet Report 2018, United Nations Environment Programme, 2018. Available at: https://livereport.protectedplanet.net/pdf/Protected_Planet_Report_2018.pdf.
- [1039] J. C. Holloway, C. Humphreys, *The business of tourism*, Sage, 2019.
- [1040] M. Scheffer, W. Brock, F. Westley, Socioeconomic mechanisms preventing optimum use of ecosystem services: an interdisciplinary theoretical analysis, *Ecosystems* 3 (2000) 451–471.
- [1041] M. Scheffer, S. Carpenter, J. A. Foley, C. Folke, B. Walker, Catastrophic shifts in ecosystems, *Nature* 413 (2001) 591–596.
- [1042] T. Klanjscek, E. B. Muller, R. M. Nisbet, Feedbacks and tipping points in organismal response to oxidative stress, *J. Theor. Biol.* 404 (2016) 361–374.
- [1043] N. Marn, M. Jusup, S. A. L. M. Kooijman, T. Klanjscek, Quantifying impacts of plastic debris on marine wildlife identifies ecological breakpoints, *Ecol. Lett.* 23 (2020) 1479–1487.

- [1044] L. C. Hamilton, M. J. Butler, Outport adaptations: Social indicators through Newfoundland's cod crisis, *Hum. Ecol. Rev.* 8 (2001) 1–11.
- [1045] L. C. Hamilton, R. L. Haedrich, C. M. Duncan, Above and below the water: social/ecological transformation in northwest Newfoundland, *Popul. Environ.* 25 (2004) 195–215.
- [1046] R. A. Myers, J. A. Hutchings, N. J. Barrowman, Why do fish stocks collapse? the example of cod in Atlantic Canada, *Ecol. Appl.* 7 (1997) 91–106.
- [1047] K. T. Frank, B. Petrie, J. A. D. Fisher, W. C. Leggett, Transient dynamics of an altered large marine ecosystem, *Nature* 477 (2011) 86–89.
- [1048] G. A. Rose, S. Rowe, Northern cod comeback, *Can. J. Fish. Aquat. Sci.* 72 (2015) 1789–1798.
- [1049] A. M. De Roos, L. Persson, Size-dependent life-history traits promote catastrophic collapses of top predators, *Proc. Natl. Acad. Sci. USA* 99 (2002) 12907–12912.
- [1050] E. King, J. Cavender-Bares, P. Balvanera, T. H. Mwampamba, S. Polasky, Trade-offs in ecosystem services and varying stakeholder preferences: evaluating conflicts, obstacles, and opportunities, *Ecol. Soc.* 20 (2015) 25.
- [1051] H. Okayama, General strategy for understanding intracellular molecular interaction cascades that elicit stimulus-invoked biological processes, *Proc. Jpn. Acad. Ser. B Phys. Biol. Sci.* 92 (2016) 372–385.
- [1052] J. A. Baross, S. E. Hoffman, Submarine hydrothermal vents and associated gradient environments as sites for the origin and evolution of life, *Orig. Life Evol. Biosph.* 15 (1985) 327–345.
- [1053] R. L. Lindeman, The trophic-dynamic aspect of ecology, *Ecology* 23 (1942) 399–417.
- [1054] R. K. Pachauri, M. R. Allen, V. R. Barros, J. Broome, W. Cramer, R. Christ, J. A. Church, L. Clarke, Q. Dahe, P. Dasgupta, et al., *Climate change 2014: synthesis report. Contribution of Working Groups*

I, II and III to the fifth assessment report of the Intergovernmental Panel on Climate Change, IPCC, 2014.

- [1055] P. W. Holland, Statistics and causal inference, *J. Am. Stat. Assoc.* 81 (1986) 945–960.
- [1056] T. A. Carleton, S. M. Hsiang, Social and economic impacts of climate, *Science* 353 (2016) aad9837.
- [1057] Y. Guo, A. Gasparrini, B. Armstrong, S. Li, B. Tawatsupa, A. Tobias, E. Lavigne, M. d. S. Z. S. Coelho, M. Leone, X. Pan, et al., Global variation in the effects of ambient temperature on mortality: a systematic evaluation, *Epidemiology* 25 (2014) 781–789.
- [1058] R. Burgess, O. Deschenes, D. Donaldson, M. Greenstone, The unequal effects of weather and climate change: Evidence from mortality in India, 2013. Available at: https://cla.umn.edu/sites/cla.umn.edu/files/burgess_deschenes_donaldson_greenstone_unequal_effects.pdf. Archived at: <https://doi.org/10.17605/OSF.IO/N7G9X>.
- [1059] S. M. Hsiang, D. Narita, Adaptation to cyclone risk: Evidence from the global cross-section, *Clim. Chang. Econ.* 3 (2012) 1250011.
- [1060] J. Anttila-Hughes, S. Hsiang, Destruction, disinvestment, and death: Economic and human losses following environmental disaster, 2013. Social Science Research Network (SSRN) Paper No. 2220501. Available at: <https://ssrn.com/abstract=2220501>.
- [1061] W. Schlenker, M. J. Roberts, Nonlinear temperature effects indicate severe damages to US crop yields under climate change, *Proc. Natl. Acad. Sci. USA* 106 (2009) 15594–15598.
- [1062] F. D. Hidalgo, S. Naidu, S. Nichter, N. Richardson, Economic determinants of land invasions, *Rev. Econ. Stat.* 92 (2010) 505–523.
- [1063] J. Graff Zivin, M. Neidell, Temperature and the allocation of time: Implications for climate change, *J. Labor Econ.* 32 (2014) 1–26.
- [1064] J. Graff Zivin, S. M. Hsiang, M. Neidell, Temperature and human capital in the short and long run, *J. Assoc. Environ. Resour. Econ.* 5 (2018) 77–105.

- [1065] M. Auffhammer, Climate adaptive response estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption using big data, 2018. National Bureau of Economic Research (NBER) Working Paper No. 24397. Available at: <https://www.nber.org/papers/w24397>.
- [1066] M. Burke, S. M. Hsiang, E. Miguel, Global non-linear effect of temperature on economic production, *Nature* 527 (2015) 235–239.
- [1067] T. Deryugina, S. M. Hsiang, Does the environment still matter? Daily temperature and income in the United States, 2014. National Bureau of Economic Research (NBER) Working Paper No. 20750. Available at: <https://www.nber.org/papers/w20750>.
- [1068] P. Zhang, O. Deschenes, K. Meng, J. Zhang, Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants, *J. Environ. Econ. Manage.* 88 (2018) 1–17.
- [1069] P. Baylis, Temperature and temperament: Evidence from Twitter, *J. Public Econ.* 184 (2020) 104161.
- [1070] M. Ranson, Crime, weather, and climate change, *J. Environ. Econ. Manage.* 67 (2014) 274–302.
- [1071] S. M. Hsiang, K. C. Meng, M. A. Cane, Civil conflicts are associated with the global climate, *Nature* 476 (2011) 438–441.
- [1072] P. Bohra-Mishra, M. Oppenheimer, S. M. Hsiang, Nonlinear permanent migration response to climatic variations but minimal response to disasters, *Proc. Natl. Acad. Sci. USA* 111 (2014) 9780–9785.
- [1073] W. Steffen, K. Richardson, J. Rockström, H. J. Schellnhuber, O. P. Dube, S. Dutreuil, T. M. Lenton, J. Lubchenco, The emergence and evolution of Earth System Science, *Nat. Rev. Earth Environ.* 1 (2020) 54–63.
- [1074] S. Seneviratne, N. Nicholls, D. Easterling, C. Goodess, S. Kanae, J. Kossin, Y. Luo, J. Marengo, K. McInnes, M. Rahimi, et al., Changes in climate extremes and their impacts on the natural physical environment, in: C. B. Field, V. Barros, T. F. Stocker, D. Qin, D. J.

- Dokken, K. L. Ebi, M. D. Mastrandrea, K. J. Mach, G.-K. Plattner, S. K. Allen, M. Tignor, P. M. Midgley (Eds.), *Managing the Risks of Extreme Events and Disasters to Advance Climate Change Adaptation. A Special Report of Working Groups I and II of the Intergovernmental Panel on Climate Change (IPCC)*, Cambridge University Press, 2012, pp. 109–230.
- [1075] T. M. Lenton, H. Held, E. Kriegler, J. W. Hall, W. Lucht, S. Rahmstorf, H. J. Schellnhuber, Tipping elements in the Earth’s climate system, *Proc. Natl. Acad. Sci. USA* 105 (2008) 1786–1793.
- [1076] M. Gladwell, *The tipping point: How little things can make a big difference*, Little, Brown, 2006.
- [1077] S. Rahmstorf, Ocean circulation and climate during the past 120,000 years, *Nature* 419 (2002) 207–214.
- [1078] V. Dakos, M. Scheffer, E. H. van Nes, V. Brovkin, V. Petoukhov, H. Held, Slowing down as an early warning signal for abrupt climate change, *Proc. Natl. Acad. Sci. USA* 105 (2008) 14308–14312.
- [1079] L. Caesar, S. Rahmstorf, A. Robinson, G. Feulner, V. Saba, Observed fingerprint of a weakening Atlantic Ocean overturning circulation, *Nature* 556 (2018) 191–196.
- [1080] C. Doyle, S. Sreenivasan, B. K. Szymanski, G. Korniss, Social consensus and tipping points with opinion inertia, *Physica A* 443 (2016) 316–323.
- [1081] D. Centola, The spread of behavior in an online social network experiment, *Science* 329 (2010) 1194–1197.
- [1082] T. M. Lenton, J. Rockström, O. Gaffney, S. Rahmstorf, K. Richardson, W. Steffen, H. J. Schellnhuber, Climate tipping points—too risky to bet against, *Nature* 575 (2019) 592–595.
- [1083] I. M. Otto, J. F. Donges, R. Cremades, A. Bhowmik, R. J. Hewitt, W. Lucht, J. Rockström, F. Allerberger, M. McCaffrey, S. S. Doe, et al., Social tipping dynamics for stabilizing Earth’s climate by 2050, *Proc. Natl. Acad. Sci. USA* 117 (2020) 2354–2365.

- [1084] S. R. Carpenter, W. A. Brock, Rising variance: a leading indicator of ecological transition, *Ecology Lett.* 9 (2006) 311–318.
- [1085] V. N. Livina, T. M. Lenton, A modified method for detecting incipient bifurcations in a dynamical system, *Geophys. Res. Lett.* 34 (2007) L03712.
- [1086] J. Prettyman, T. Kuna, V. Livina, A novel scaling indicator of early warning signals helps anticipate tropical cyclones, *EPL (Europhy. Lett.)* 121 (2018) 10002.
- [1087] R. Wang, J. A. Dearing, P. G. Langdon, E. Zhang, X. Yang, V. Dakos, M. Scheffer, Flickering gives early warning signals of a critical transition to a eutrophic lake state, *Nature* 492 (2012) 419–422.
- [1088] V. Guttal, C. Jayaprakash, Changing skewness: an early warning signal of regime shifts in ecosystems, *Ecol. Lett.* 11 (2008) 450–460.
- [1089] R. Biggs, S. R. Carpenter, W. A. Brock, Turning back from the brink: detecting an impending regime shift in time to avert it, *Proc. Natl. Acad. Sci. USA* 106 (2009) 826–831.
- [1090] V. Dakos, S. R. Carpenter, W. A. Brock, A. M. Ellison, V. Guttal, A. R. Ives, S. Kefi, V. Livina, D. A. Seekell, E. H. van Nes, et al., Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data, *PLOS ONE* 7 (2012) e41010.
- [1091] L. F. Richardson, *Weather prediction by numerical process*, Cambridge university press, 1922.
- [1092] G. S. Callendar, The artificial production of carbon dioxide and its influence on temperature, *Q. J. R. Meteorol. Soc.* 64 (1938) 223–240.
- [1093] S. Manabe, R. T. Wetherald, Thermal equilibrium of the atmosphere with a given distribution of relative humidity, *J. Atmos. Sci.* 24 (1967) 241–259.
- [1094] Working Group 1 of the IPCC, Policymakers summary, in: J. T. Houghton, G. J. Jenkins, J. J. Ephraums (Eds.), *Climate Change: The IPCC Scientific Assessment*, Cambridge University Press, 1990, pp. vii–xxxiv.

- [1095] J. Rockström, W. Steffen, K. Noone, Å. Persson, F. S. Chapin, E. F. Lambin, T. M. Lenton, M. Scheffer, C. Folke, H. J. Schellnhuber, et al., A safe operating space for humanity, *Nature* 461 (2009) 472–475.
- [1096] I. Medhaug, M. B. Stolpe, E. M. Fischer, R. Knutti, Reconciling controversies about the ‘global warming hiatus’, *Nature* 545 (2017) 41–47.
- [1097] J. D. Neelin, *Climate change and climate modeling*, Cambridge University Press, 2010.
- [1098] S. E. Zebiak, M. A. Cane, A model of El Niño–southern oscillation, *Mon. Weather Rev.* 115 (1987) 2262–2278.
- [1099] E. A. Parson, K. Fisher-Vanden, Integrated assessment models of global climate change, *Annu. Rev. Energ. Environ.* 22 (1997) 589–628.
- [1100] D. P. van Vuuren, J. Lowe, E. Stehfest, L. Gohar, A. F. Hof, C. Hope, R. Warren, M. Meinshausen, G.-K. Plattner, How well do integrated assessment models simulate climate change?, *Clim. Change* 104 (2011) 255–285.
- [1101] C. Hope, J. Anderson, P. Wenman, Policy analysis of the greenhouse effect: an application of the PAGE model, *Energy Policy* 21 (1993) 327–338.
- [1102] P. Holme, J. C. Rocha, *Networks of climate change: Connecting causes and consequences*, 2021. e-print arXiv:2105.12537.
- [1103] A. A. Tsonis, P. J. Roebber, The architecture of the climate network, *Physica A* 333 (2004) 497–504.
- [1104] A. A. Tsonis, K. L. Swanson, P. J. Roebber, What do networks have to do with climate?, *Bull. Am. Meteorol. Soc.* 87 (2006) 585–596.
- [1105] A. A. Tsonis, K. L. Swanson, Topology and predictability of El Niño and La Niña networks, *Phys. Rev. Lett.* 100 (2008) 228502.

- [1106] K. Yamasaki, A. Gozolchiani, S. Havlin, Climate networks around the globe are significantly affected by El Niño, *Phys. Rev. Lett.* 100 (2008) 228501.
- [1107] A. Gozolchiani, K. Yamasaki, O. Gazit, S. Havlin, Pattern of climate network blinking links follows El Niño events, *EPL (Europhys. Lett.)* 83 (2008) 28005.
- [1108] A. Gozolchiani, S. Havlin, K. Yamasaki, Emergence of El Niño as an autonomous component in the climate network, *Phys. Rev. Lett.* 107 (2011) 148501.
- [1109] J. Ludescher, A. Gozolchiani, M. I. Bogachev, A. Bunde, S. Havlin, H. J. Schellnhuber, Improved El Niño forecasting by cooperativity detection, *Proc. Natl. Acad. Sci. USA* 110 (2013) 11742–11745.
- [1110] J. Meng, J. Fan, Y. Ashkenazy, S. Havlin, Percolation framework to describe El Niño conditions, *Chaos* 27 (2017) 035807.
- [1111] J. Fan, J. Meng, Y. Ashkenazy, S. Havlin, H. J. Schellnhuber, Network analysis reveals strongly localized impacts of El Niño, *Proc. Natl. Acad. Sci. USA* 114 (2017) 7543–7548.
- [1112] J. Meng, J. Fan, Y. Ashkenazy, A. Bunde, S. Havlin, Forecasting the magnitude and onset of El Niño based on climate network, *New J. Phys.* 20 (2018) 043036.
- [1113] N. Boers, B. Bookhagen, H. M. Barbosa, N. Marwan, J. Kurths, J. Marengo, Prediction of extreme floods in the eastern Central Andes based on a complex networks approach, *Nat. Commun.* 5 (2014) 5199.
- [1114] N. Boers, B. Goswami, A. Rheinwalt, B. Bookhagen, B. Hoskins, J. Kurths, Complex networks reveal global pattern of extreme-rainfall teleconnections, *Nature* 566 (2019) 373–377.
- [1115] V. Stolbova, E. Surovyatkina, B. Bookhagen, J. Kurths, Tipping elements of the Indian monsoon: Prediction of onset and withdrawal, *Geophys. Res. Lett.* 43 (2016) 3982–3990.

- [1116] M. Van Der Mheen, H. A. Dijkstra, A. Gozolchiani, M. Den Toom, Q. Feng, J. Kurths, E. Hernandez-Garcia, Interaction network based early warning indicators for the Atlantic MOC collapse, *Geophys. Res. Lett.* 40 (2013) 2714–2719.
- [1117] Q. Y. Feng, J. P. Viebahn, H. A. Dijkstra, Deep ocean early warning signals of an Atlantic MOC collapse, *Geophys. Res. Lett.* 41 (2014) 6009–6015.
- [1118] Q. Y. Feng, H. Dijkstra, Are North Atlantic multidecadal SST anomalies westward propagating?, *Geophys. Res. Lett.* 41 (2014) 541–546.
- [1119] D. Zhou, A. Gozolchiani, Y. Ashkenazy, S. Havlin, Teleconnection paths via climate network direct link detection, *Phys. Rev. Lett.* 115 (2015) 268501.
- [1120] N. Ying, D. Zhou, Z. Han, Q. Chen, Q. Ye, Z. Xue, Rossby waves detection in the CO2 and temperature multilayer climate network, *Geophys. Res. Lett.* 47 (2020) e2019GL086507.
- [1121] J. F. Donges, J. Heitzig, B. Beronov, M. Wiedermann, J. Runge, Q. Y. Feng, L. Tupikina, V. Stolbova, R. V. Donner, N. Marwan, et al., Unified functional network and nonlinear time series analysis for complex systems science: The pyunicorn package, *Chaos* 25 (2015) 113101.
- [1122] J. Fan, J. Meng, Y. Ashkenazy, S. Havlin, H. J. Schellnhuber, Climate network percolation reveals the expansion and weakening of the tropical component under global warming, *Proc. Natl. Acad. Sci. USA* 115 (2018) E12128–E12134.
- [1123] R. Q. Quiroga, T. Kreuz, P. Grassberger, Event synchronization: a simple and fast method to measure synchronicity and time delay patterns, *Phys. Rev. E* 66 (2002) 041904.
- [1124] R. V. Donner, Y. Zou, J. F. Donges, N. Marwan, J. Kurths, Recurrence networks—a novel paradigm for nonlinear time series analysis, *New J. Phys.* 12 (2010) 033025.

- [1125] H. A. Dijkstra, *Nonlinear physical oceanography: a dynamical systems approach to the large scale ocean circulation and El Niño*, Springer Science & Business Media, 2005.
- [1126] A. J. Clarke, *An introduction to the dynamics of El Niño and the Southern Oscillation*, Elsevier, 2008.
- [1127] W. Duan, C. Wei, The ‘spring predictability barrier’ for ENSO predictions and its possible mechanism: results from a fully coupled model, *Int. J. Climatol.* 33 (2013) 1280–1292.
- [1128] J. Ludescher, A. Gozolchiani, M. I. Bogachev, A. Bunde, S. Havlin, H. J. Schellnhuber, Very early warning of next El Niño, *Proc. Natl. Acad. Sci. USA* 111 (2014) 2064–2066.
- [1129] B. P. Kirtman, The COLA anomaly coupled model: Ensemble ENSO prediction, *Mon. Weather Rev.* 131 (2003) 2324–2341.
- [1130] D. Chen, M. A. Cane, A. Kaplan, S. E. Zebiak, D. Huang, Predictability of El Niño over the past 148 years, *Nature* 428 (2004) 733–736.
- [1131] N. K. Larkin, D. Harrison, On the definition of El Niño and associated seasonal average US weather anomalies, *Geophys. Res. Lett.* 32 (2005) L13705.
- [1132] H.-Y. Kao, J.-Y. Yu, Contrasting eastern-Pacific and central-Pacific types of ENSO, *J. Clim.* 22 (2009) 615–632.
- [1133] K. Ashok, S. K. Behera, S. A. Rao, H. Weng, T. Yamagata, El Niño Modoki and its possible teleconnection, *J. Geophys. Res. Oceans* 112 (2007) C11007.
- [1134] Z. Lu, N. Yuan, L. Chen, Z. Gong, On the impacts of El Niño events: A new monitoring approach using complex network analysis, *Geophys. Res. Lett.* 47 (2020) e2019GL086533.
- [1135] M. K. Roxy, S. Ghosh, A. Pathak, R. Athulya, M. Mujumdar, R. Murugudde, P. Terray, M. Rajeevan, A threefold rise in widespread extreme rain events over central India, *Nat. Commun.* 8 (2017) 1–11.

- [1136] R. Son, S.-Y. S. Wang, W.-L. Tseng, C. W. B. Schuler, E. Becker, J.-H. Yoon, Climate diagnostics of the extreme floods in Peru during early 2017, *Clim. Dyn.* 54 (2020) 935–945.
- [1137] E. M. Fischer, R. Knutti, Observed heavy precipitation increase confirms theory and early models, *Nat. Clim. Change* 6 (2016) 986–991.
- [1138] A. Borodina, E. M. Fischer, R. Knutti, Models are likely to underestimate increase in heavy rainfall in the extratropical regions with high rainfall intensity, *Geophys. Res. Lett.* 44 (2017) 7401–7409.
- [1139] A. Guterres, Stressing air pollution kills 7 million people annually, Secretary-General urges governments to build green economy, in message for World Environment Day, 2019. Press release SG/SM/19607-ENV/DEV/1957-OBV/1887. Available at: <https://www.un.org/press/en/2019/sgsm19607.doc.htm>.
- [1140] Y. Zhang, J. Fan, X. Chen, Y. Ashkenazy, S. Havlin, Significant impact of Rossby waves on air pollution detected by network analysis, *Geophys. Res. Lett.* 46 (2019) 12476–12485.
- [1141] M. Scheffer, *Critical transitions in nature and society*, Princeton University Press, 2009.
- [1142] H. E. Stanley, *Introduction to phase transitions and critical phenomena*, Oxford University Press, 1971.
- [1143] O. Peters, J. D. Neelin, Critical phenomena in atmospheric precipitation, *Nat. Phys.* 2 (2006) 393–396.
- [1144] V. Privman, *Finite size scaling and numerical simulation of statistical systems*, World Scientific, 1990.
- [1145] I. M. Held, B. J. Soden, Robust responses of the hydrological cycle to global warming, *J. Clim.* 19 (2006) 5686–5699.
- [1146] Q. Fu, C. M. Johanson, J. M. Wallace, T. Reichler, Enhanced mid-latitude tropospheric warming in satellite measurements, *Science* 312 (2006) 1179–1179.
- [1147] J. Lu, G. A. Vecchi, T. Reichler, Expansion of the Hadley cell under global warming, *Geophys. Res. Lett.* 34 (2007) L06805.

- [1148] K.-H. Seo, D. M. Frierson, J.-H. Son, A mechanism for future changes in Hadley circulation strength in CMIP5 climate change simulations, *Geophys. Res. Lett.* 41 (2014) 5251–5258.
- [1149] C. M. Mitas, A. Clement, Has the Hadley cell been strengthening in recent decades?, *Geophys. Res. Lett.* 32 (2005) L03809.
- [1150] D. Stauffer, A. Aharony, *Introduction to percolation theory*, Taylor & Francis, 2003.
- [1151] J. Fan, J. Meng, Y. Liu, A. A. Saberi, J. Kurths, J. Nagler, Universal gap scaling in percolation, *Nat. Phys.* 16 (2020) 455–461.
- [1152] J. P. Kossin, K. A. Emanuel, G. A. Vecchi, The poleward migration of the location of tropical cyclone maximum intensity, *Nature* 509 (2014) 349–352.
- [1153] H. Dijkstra, P. Petersik, E. Hernandez-Garcia, C. Lopez, The application of machine learning techniques to improve El Niño prediction skill, *Front. Phys.* 7 (2019) 153.
- [1154] W. Luo, M. Steptoe, Z. Chang, R. Link, L. Clarke, R. Maciejewski, Impact of spatial scales on the intercomparison of climate scenarios, *IEEE Comput. Graph. Appl.* 37 (2017) 40–49.
- [1155] J. Berezovsky, The structure of musical harmony as an ordered phase of sound: A statistical mechanics approach to music theory, *Sci. Adv.* 5 (2019) eaav8490.
- [1156] D. Kim, S.-W. Son, H. Jeong, Large-scale quantitative analysis of painting arts, *Sci. Rep.* 4 (2014) 7370.
- [1157] K. Laws, *Physics and the Art of Dance: Understanding Movement*, Oxford University Press, 2002.
- [1158] Y. Gandica, J. S. Lansing, N. N. Chung, S. Thurner, L. Y. Chew, Bali’s ancient rice terraces: A Hamiltonian approach, 2021. e-print arXiv:2103.04466.
- [1159] Y.-Y. Ahn, S. E. Ahnert, J. P. Bagrow, A.-L. Barabási, Flavor network and the principles of food pairing, *Sci. Rep.* 1 (2011) 196.

- [1160] S. Lee, V. C. Ramenzoni, P. Holme, Emergence of collective memories, PLOS ONE 5 (2010) e12522.
- [1161] D. Braha, Global civil unrest: contagion, self-organization, and prediction, PLOS ONE 7 (2012) e48596.