

RESEARCH ARTICLE

A Design Pattern for Decentralised Decision Making

Andreagiovanni Reina¹✉, Gabriele Valentini¹, Cristian Fernández-Oto², Marco Dorigo¹, Vito Trianni³✉*

1 IRIDIA, Université Libre de Bruxelles, Brussels, Belgium, **2** Faculté des Sciences, Université Libre de Bruxelles, Brussels, Belgium, **3** ISTC, Consiglio Nazionale delle Ricerche, Rome, Italy

✉ These authors contributed equally to this work.

* vito.trianni@istc.cnr.it



CrossMark
click for updates

OPEN ACCESS

Citation: Reina A, Valentini G, Fernández-Oto C, Dorigo M, Trianni V (2015) A Design Pattern for Decentralised Decision Making. PLoS ONE 10(10): e0140950. doi:10.1371/journal.pone.0140950

Editor: Zhen Wang, Kyushu University, JAPAN

Received: June 26, 2015

Accepted: October 2, 2015

Published: October 23, 2015

Copyright: © 2015 Reina et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: "All data have been obtained through a multi agent simulation software based on the MASON library (<https://cs.gmu.edu/~eclab/projects/mason>). The code for running the multi-agent simulations of the case studies is available on Github at the following address: <https://github.com/joefresna/PlosONE-Reina2015>.

Funding: The research presented in this paper has been supported by the European Research Council (ERC) Advanced Grant "E-SWARM: Engineering Swarm Intelligence Systems" awarded to MD (contract 246939, <http://www.e-swarm.org>) and by the Marie Skłodowska-Curie Career Integration Grant "DICE: Distributed Cognition Engineering" awarded to VT (contract 631297, <http://laral.istc.cnr.it/dice->

Abstract

The engineering of large-scale decentralised systems requires sound methodologies to guarantee the attainment of the desired macroscopic system-level behaviour given the microscopic individual-level implementation. While a general-purpose methodology is currently out of reach, specific solutions can be given to broad classes of problems by means of well-conceived *design patterns*. We propose a design pattern for collective decision making grounded on experimental/theoretical studies of the nest-site selection behaviour observed in honeybee swarms (*Apis mellifera*). The way in which honeybee swarms arrive at consensus is fairly well-understood at the macroscopic level. We provide formal guidelines for the microscopic implementation of collective decisions to quantitatively match the macroscopic predictions. We discuss implementation strategies based on both homogeneous and heterogeneous multiagent systems, and we provide means to deal with spatial and topological factors that have a bearing on the micro-macro link. Finally, we exploit the design pattern in two case studies that showcase the viability of the approach. Besides engineering, such a design pattern can prove useful for a deeper understanding of decision making in natural systems thanks to the inclusion of individual heterogeneities and spatial factors, which are often disregarded in theoretical modelling.

Introduction

Large-scale decentralised systems are becoming the more ubiquitous the more our lives become connected, and the problem of understanding and controlling such complex systems is currently approached from a variety of perspectives [1–3]. A particularly interesting viewpoint on decentralised systems is the one of Cognitive Science, which looks at their ability of collecting, transforming and propagating information within the system while interacting with the external world [4, 5]. However, providing desired cognitive processing abilities to large-scale distributed systems—from swarms of robots [6] to cognitive radio networks [7] and cyber-physical systems [8]—is a complex problem that has no general solution available to date. This is mainly due to difficulties in treating at the same time several complexity factors (e.g., heterogeneities

project). MD acknowledges support from the Belgian "Fonds de la Recherche Scientifique - FNRS" (<http://www.frs-fnrs.be/>). CFO acknowledges support from Becas Chile (http://www.mineduc.cl/index.php?id_portal=60).

Competing Interests: The authors have declared that no competing interests exist.

in the interaction topology or in the individual behaviour). Additionally, each application faces domain-specific challenges. For instance, in swarm robotics, any design methodology needs to deal with the inherent spatial factors [9–12].

In lieu of a universal design methodology, general solutions can be provided for specific classes of problems exploiting the concept of design patterns well known in software engineering [13]. Design patterns provide formal guidelines to deal with recurring problems in a specific field. For distributed systems, design patterns prescribe the individual-level microscopic behaviour required to obtain desired system-level macroscopic properties [14–16]. According to the literature, design patterns are defined by the following attributes: name, context, design rationale, problem and solution (see S1 Text). Each attribute identifies and formally describes one component of the proposed design method. For instance, the context determines the domain of applicability of the design pattern, while the design rationale explains the what and the how, that is, the origins and working principles. Additionally, one or more case studies illustrate the application of the proposed solution to a specific problem instance.

Among the classes of cognitive processes that can be addressed through well-conceived design patterns, decentralised decision making represents a fundamental ability in several contexts and application domains [17–20]. In this paper, we propose a design pattern for decentralised decision making based on the nest-site selection behaviour of honeybee swarms [21–23]. Previous experimental and theoretical studies have demonstrated near-optimal speed-accuracy tradeoffs in the selection of the most profitable option among a set of alternative nesting sites by honeybees [21, 22]. Most importantly, inhibitory signals among bees provide an adaptive mechanism to quickly break deadlocks and tune the decision dynamics according to the perceived quality of the discovered options [22, 23]. The above properties of the nest-site selection process are relevant for many practical decision-making scenarios in decentralised systems, and justify its choice in this study.

Starting from the macroscopic description of the nest-site selection dynamics [22, 23], we derive the exact relationship between microscopic and macroscopic models—also including finite-size effects—for the general case of a best-of- n decision problem. The inter-related models represent the core of the design pattern, which is completed by formal guidelines for the implementation of collective decisions in multiagent systems. We provide guidelines for implementation by means of either homogenous or heterogeneous agents, as well as guidelines for the inclusion of spatial and topological factors that have a bearing in determining the microscopic interaction patterns [9]. We illustrate the application of the design pattern in two case studies, covering decentralised decisions by (i) static agents interacting in a fully-connected network, and (ii) mobile agents involved in a search and exploitation task.

Collective Decisions through Cross Inhibition

We consider a best-of- n decision problem, that is, the choice of the best option, or any of the equal-best options, among n different alternatives. Each option $i \in \{1, \dots, n\}$ is characterised by a quality $v_i \in [v_m, v_M]$. We study decision making for a population \mathcal{A} of N agents, where each agent $a_g, g \in \{1, \dots, N\}$ is either committed to one of the available options i and belongs to the sub-population \mathcal{A}_i (size N_i and fraction $\Psi_i = N_i/N$) or is uncommitted and belongs to sub-population \mathcal{U} (size N_U and fraction Ψ_U). Agents can obtain a noisy estimate \hat{v}_i of the quality associated to option i . At the macroscopic level, a decision is taken as soon as the entire population or a large fraction Ψ_q (hereafter referred to as *quorum*) is committed to a single option.

In this section, we first present a macroscopic model that accounts for the system mean-field dynamics, and that describes the decision process at the population level without

considering the actions of each individual agent and their inherent stochasticity. Then, we introduce a microscopic model in the form of a probabilistic finite state machine (PFSM) to describe the behaviour of the individual agent that changes its commitment state in response to probabilistic events. This description level is the closest to the agent implementation and allows us to directly link the control parameters of the agent behaviour to the macroscopic parameters. Finally, we provide the link between macro and micro levels through equations that convert ODE transition rates into the PFSM probabilities.

Macroscopic model

According to the model proposed for honeybee nest-site selection [22], four concurrent processes determine the distribution of agents across populations: (i) uncommitted agents spontaneously *discover* option i at rate γ_i (ii) agents committed to option i spontaneously *abandon* commitment at rate α_i ; (iii) agents committed to option i *recruit* uncommitted agents at rate ρ_i ; and (iv) agents committed to option $j \neq i$ *inhibit* agents committed to option i at rate σ_j (cross-inhibition). A committed agent that receives an inhibitory signal stops recruiting (hence the expression *stop signal*) and reverts to the uncommitted state, becoming available for discovery or recruitment by other agents. Such cross-inhibition mechanism allows the swarm to tune the decision speed and break decision deadlocks in case of equal-best options [22].

The mean-field macroscopic dynamics are well described by an n -dimensional ODE system, which extends the previously studied binary version [22, 23]:

$$\begin{cases} \dot{\Psi}_i = \gamma_i \Psi_U - \alpha_i \Psi_i + \rho_i \Psi_i \Psi_U - \sum_{j \neq i} \sigma_j \Psi_i \Psi_j \\ \Psi_U = 1 - \sum_i \Psi_i \end{cases} \quad (1)$$

where $i \in \{1, \dots, n\}$. Here, the variation of Ψ_i is determined by the four processes described above, and a mass conservation equation constrains the dynamics within feasible bounds. The transition rates ($\gamma_i, \alpha_i, \rho_i$ and σ_i) are functions of the quality v_i :

$$\alpha_i = f_\alpha(v_i), \quad \gamma_i = f_\gamma(v_i), \quad \rho_i = f_\rho(v_i), \quad \sigma_i = f_\sigma(v_i) \quad (2)$$

The relations between option quality and transition rates determine the macroscopic dynamics [23].

Microscopic model

The behaviour of individual agents is represented by the probabilistic finite state machine (PFSM) shown in Fig 1A that describes the *commitment dynamics*. An agent can be either uncommitted (state C_U) or committed to option i (state C_i), and probabilistically changes state every τ seconds according to two types of transitions: spontaneous and interactive. Spontaneous transitions model the discovery of option i with probability $\mathcal{P}_\gamma(v_i)$ and the abandonment of commitment to option i with probability $\mathcal{P}_\alpha(v_i)$. Interactive transitions model the recruitment and cross-inhibition processes resulting from the interaction between agents belonging to different populations. We refer to the probability of any agent interacting with an agent committed to option i as \mathcal{P}_{Ψ_i} . We assume a well-mixed system, so that $\mathcal{P}_{\Psi_i} = N_i/N$. Recruitment for option i is modelled by a transition from C_U to C_i with overall probability $\mathcal{P}_{\Psi_i} \mathcal{P}_\rho(v_i)$. Cross-inhibition of an agent committed to option i is instead modelled as the cumulative effect of the interaction with agents committed to a different option, with overall probability $\sum_{j \neq i} \mathcal{P}_{\Psi_j} \mathcal{P}_\sigma(v_j)$ (see S1 Text for details).




Fig 1. Probabilistic Finite State Machines (PFSMs) describing the microscopic behaviour of an agent in average. Here, the notation \mathcal{P}_i is a shorthand for $\mathcal{P}_i(v_i)$ (A) PFSM describing the basic commitment dynamics. Spontaneous transitions are represented by solid lines, interactive transitions by dashed lines. For a complete version with $n + 1$ states, see panel A in S1 Fig (B) PFSM describing the dynamics of activity change as agents switch from latent to interactive states. Latent states are indicated in grey, while changes in the activity state are represented by dash-dotted arrows. (C) PFSM describing the coupled dynamics of commitment and activity change. For a complete version with $2(n + 1)$ states, see panel B in S1 Fig.

doi:10.1371/journal.pone.0140950.g001

Micro-macro link

The PFSM of Fig 1A represents the average agent behaviour (see also S1 Fig) [24]. Its actual implementation requires choosing the way in which transitions are executed in relation to the limited information available to the individual agent. For instance, the estimation of the population-size dependent probability \mathcal{P}_{Ψ_i} by individual agents requires some sampling of the current population size. The implementation is rather straightforward, and is detailed in S1 Text. Here, we focus on the other transition probabilities ($\mathcal{P}_\gamma, \mathcal{P}_\alpha, \mathcal{P}_\rho, \mathcal{P}_\sigma$), and their correspondence with the macroscopic transition rates to obtain the desired dynamics determined by Eq (2).

We propose two strategies based either on a *homogeneous* or on a *heterogeneous* implementation. In the homogeneous case, all agents compute their transition probabilities in the same way as a function of the estimated quality \hat{v}_i . In this case, it is possible to establish a direct correspondence between micro and macro parameters:

$$P_{\lambda,g}(\hat{v}_i) = \lambda_i \tau = f_\lambda(\hat{v}_i) \tau, \quad \begin{aligned} i &\in \{1, \dots, n\} \\ \lambda &\in \{\gamma, \alpha, \rho, \sigma\} \end{aligned} \quad (3)$$

where $P_{\lambda,g}$ represents the actual probability for the agent a_g to undergo the transition λ . Hence, a principled choice of the individual transition probabilities can be made to obtain the desired macroscopic dynamics as defined by Eq (2). The derivation of the micro-macro link is detailed in S1 Text.

In the heterogeneous case, each agent computes its own transition probabilities differently from other agents. We propose a simple response threshold scheme [25], so that agent a_g follows a transition with a fixed probability if the (estimated) option quality \hat{v}_i exceeds a given threshold δ_g :

$$P_{\lambda,g}(\hat{v}_i) = \begin{cases} \mathcal{P}_{\lambda\uparrow} & \hat{v}_i > \delta_g \\ \mathcal{P}_{\lambda\downarrow} & \hat{v}_i \leq \delta_g \end{cases}, \quad \begin{aligned} i &\in \{1, \dots, n\} \\ \lambda &\in \{\gamma, \alpha, \rho, \sigma\} \end{aligned} \quad (4)$$

where $\mathcal{P}_{\lambda\uparrow}$ and $\mathcal{P}_{\lambda\downarrow}$ are tuneable parameters, and the value δ_g is drawn for each agent a_g from a probability distribution \mathcal{D}_λ over the range $[v_m, v_M]$. With this implementation, it is possible to establish a relationship between microscopic and macroscopic parameters through the

cumulative distribution function $F_{\mathcal{D}_\lambda}$ of \mathcal{D}_λ :

$$F_{\mathcal{D}_\lambda} = \frac{\lambda\tau - \mathcal{P}_{\lambda\downarrow}}{\mathcal{P}_{\lambda\uparrow} - \mathcal{P}_{\lambda\downarrow}}, \quad \lambda \in \{\gamma, \alpha, \rho, \sigma\} \tag{5}$$

For $F_{\mathcal{D}_\lambda}$ to be a cumulative distribution function, the relationship between quality and macroscopic transition rate expressed in Eq (2) must be monotonic in v —either increasing or decreasing. As a consequence, the step function Eq (4) can be determined by:

$$\begin{aligned} \mathcal{P}_{\lambda\uparrow} &= f_\lambda(v_M)\tau \\ \mathcal{P}_{\lambda\downarrow} &= f_\lambda(v_m)\tau \end{aligned} \quad \lambda \in \{\gamma, \alpha, \rho, \sigma\} \tag{6}$$

The micro-macro correspondence given by Eqs (5) and (6) holds when each agent a_g re-samples the threshold δ_g from \mathcal{D}_λ at every decision step (see also S1 Text). An approximation with fixed thresholds is valid for recruitment and cross-inhibition, because re-sampling is ensured by changing partner in each different interaction (as shown in case study I-A below). Fixed thresholds for spontaneous transitions should not be used, unless the macroscopic dynamics are dominated by recruitment and cross-inhibition (see case study I-B).

For both strategies—homogeneous and heterogeneous—the derivation of the relationship between microscopic and macroscopic description levels passes through the introduction of a finite-size stochastic macroscopic model determined by the system Master equation, which allows us to link the microscopic stochastic description with the macroscopic mean-field dynamics (see S1 Text for a detailed description).

Latent and interactive agents

In a practical application scenario, agents might not be able to interact with neighbours every τ seconds. For instance, an agent might be busy estimating the quality of a discovered option, or spatial/topological factors might prevent frequent interactions. Agents unable to interact are latent, as opposed to interactive ones. We model changes in this activity state (i.e., the *activity dynamics*) by considering that an agent becomes latent with probability \mathcal{P}_L , and returns interactive with probability \mathcal{P}_I . The switch in activity state is represented by the PFSM in Fig 1B, which predicts that a fraction of $\eta_I = \mathcal{P}_I/(\mathcal{P}_I + \mathcal{P}_L)$ agents can be found asymptotically in the interactive state. Similarly, $\eta_L = \mathcal{P}_L/(\mathcal{P}_I + \mathcal{P}_L)$ represents the fraction of agents in the latent state. The activity switch is possible for both committed and uncommitted agents, leading to a PFSM description with $2(n + 1)$ states (see S1 Text). Recruitment and cross-inhibition are available only to interactive agents, while discovery and abandonment are available also to latent agents, depending on the particular application. As an example, Fig 1C reports a PFSM that correctly preserves the micro-macro link, and is exploited for the search and exploration task discussed below. Given the microscopic description, it is possible to derive the correspondence between micro and macro parameters by dividing the macroscopic transition rates by η_I or η_L in Eqs (3), (5) and (6). For the specific case shown in Fig 1C, recruitment and cross-inhibition are modulated by η_B , discovery and abandonment by η_L (see S1 Text).

Results

Following the implementation guidelines described above, we present here two case studies that showcase the usage of the proposed design pattern. Case study I concerns decentralised decisions-making by static agents interacting on a fully-connected network, and is divided in two parts, I-A and I-B. Case study II concerns collective decisions by mobile agents in the context of a search and exploitation task. To ease the discussion and simplify the visualisation, we

present here a binary decision problem (options A and B with quality v_A and v_B), and we provide additional results for the best-of- n scenario in the online supplementary material. To quantify the agreement between macroscopic models and microscopic implementation, we look at the system performance through a set of metrics detailed below, and we compare the process dynamics at different abstraction levels.

Metrics

Different metrics are used in the literature to evaluate the results of a decision-making process, which are linked to the correctness of the response, the coherence across multiple repetitions, as well as the speed of the process. Whenever time is required to gather sufficient information, decision making gives rise to speed-accuracy tradeoffs—a very common phenomenon in biological systems [26]. In collective decision making, a high coherence of the group is also important, as it can minimise the costs for conflicting choices by individuals. As maintaining coherence is a time-consuming process that requires to spread information widely within the group, speed-cohesion tradeoffs may also appear [27].

In decentralised systems, accuracy can be defined as the proportion of the group that is committed to the best option, or to any of the equal-best options. Conversely, cohesion measures the ability of the group to be committed to the same option, notwithstanding its quality [27]. Therefore, one can simultaneously have low accuracy and high cohesion, for instance if 10% of the group chooses a high-quality option and the remaining 90% goes for a low-quality one. In the context of engineering artificial systems, it is important to quantify both the aspects of decision accuracy and group cohesion, and contrast them with the time required to arrive at a decision.

To this purpose, we first introduce the **resolution** \mathcal{R} , which refers to the ability to discriminate between different-quality options, and is related to the normalised quality difference between any two options A and B: $\mathcal{R} = |v_A - v_B| / \max(v_A, v_B)$. By requiring a target resolution, the designer can focus on maximising accuracy only in those portions of the problem space that are above resolution, and just require that any solution is chosen below the given resolution threshold. Note that resolution is normalised so that the minimum quality difference that can be detected is proportional to the quality magnitude, in analogy to many biological decision-making processes following the Weber's law.

Then, we consider the **effectivity** \mathcal{E} as the ability of the group to take a decision within the maximum execution time T . The effectivity is measured as the fraction of runs that reach the quorum Ψ_q within the given time limit. Effectivity is related to the coherence of decision making, as it measures the ability to take a decision (i.e., reach the predefined quorum) within the maximum allotted time, notwithstanding the quality of the chosen option. By requiring a minimum effectivity threshold, the designer can impose that the system reaches a coherent state within a maximum time T .

Having defined resolution and effectivity, we introduce the main performance metrics we take into account. The **success rate** \mathcal{S} corresponds to the fraction of effective runs resulting in a correct decision—i.e., the quorum is reached for the best option, or any of the equal-best options—when starting from a fully-uncommitted population. The success rate is defined in analogy to the exit probability in stochastic processes, and is related to the accuracy of decision making, because the quality of the chosen option is taken into account. Note that, by looking at the effective runs only—e.g., runs with effectivity larger than the given threshold—we limit ourselves only to high-coherence results. The **convergence time** \mathcal{C} is the average time required to reach the quorum Ψ_q computed over all effective runs, and is defined in analogy to the exit time of stochastic processes. This metric actually corresponds to the speed of the decision-

making process, and can be exploited together with the success rate to select the most convenient solution that optimises the speed-accuracy/cohesion tradeoff.

Case study I: Collective decisions on a fully-connected network

The first case study illustrates the implementation of decentralised decision making for a multiagent system in which each agent can potentially interact with any other agent (see [Methods](#)). We present two parameterisations as case study I-A and I-B.

In Case study I-A, we study consensus decisions, that is, we design a system in which the desired outcome is complete convergence of the group towards the choice of one or the other option. To this end, we set the decision quorum $\Psi_q = 1$ and we require that a decision is taken within $T = 400s$. Here, we also assume that option quality varies in $v \in [0, 1]$.

The first step towards implementation is the definition of the macroscopic parameterisation and its relationship with the option quality. The analysis of the macroscopic dynamics from [Eq \(1\)](#) reveals that consensus can be achieved only when abandonment is null ($\alpha_A = \alpha_B = 0$). We arbitrarily choose a constant cross-inhibition rate $\sigma_A = \sigma_B = \bar{\sigma}$, which is sufficient for determining a collective decision [\[23\]](#). The value $\bar{\sigma}$ can be tuned to determine the time scale of the process: the higher the rate, the quicker the convergence dynamics. Here, we choose $\bar{\sigma} = 1$. In these conditions, the model predicts two equilibrium points corresponding to consensus decision for either of the two options, but their stability may vary depending on the relative strength of discovery and recruitment. Assuming $v_A \geq v_B$, the [Model \(1\)](#) predicts that the equilibrium at consensus for A is always stable, while consensus for B is stable only when $\gamma_A < \gamma_B + \rho_B$ (see [S2 Text](#)). Thanks to this result, an informed choice can be made about the macroscopic parameterisation and the relation with option quality: $\gamma_i = f_\gamma(v_i)$, $\rho_i = f_\rho(v_i)$, $i \in \{A, B\}$. In particular, assuming a target resolution $R = 0.15$, we can minimise the chances of a wrong decision by designing the system to have a single stable equilibrium for the best option in any decision problem characterised by above-resolution quality differences. We select linear functions that link macroscopic transition rates to the quality:

$$f_\gamma(v_i) = k v_i, \quad f_\rho(v_i) = h v_i \tag{7}$$

where k and h are tuneable parameters. Next, we compute the constraint on the above functions to satisfy our design choice: $k > h(1 - R)/R$ (see [S2 Text](#)). Finally, we choose values that comply with the prescribed bounds: $h = 0.1$ and $k = 0.6$.

The second step towards implementation is the analysis of the system performance in the complete decision space for varying system size N . This can be studied numerically by approximating the finite-size macroscopic dynamics using the Gillespie algorithm [\[28\]](#). Finally, the multiagent system can be deployed following the prescriptions of the design pattern and choosing a convenient implementation strategy. In the homogeneous case, all agents determine P_γ and P_ρ in the same way according to [Eqs \(3\) and \(7\)](#). Conversely, in the heterogeneous case transition probabilities are determined by the step function of [Eq \(4\)](#), and vary from agent to agent with thresholds randomly sampled from the distribution determined by [Eqs \(5\), \(6\) and \(7\)](#). In this case study, we use fixed thresholds and we therefore limit the heterogeneous implementation to the recruitment probability P_ρ , while we keep the discovery probability P_γ homogeneous across agents. Finally, we let agents update their state every $\tau = 0.2 s$ following the design pattern guidelines (see [S1 Text](#)).

The performance of the multiagent system for both homogenous and heterogeneous implementations is compared to the macroscopic Gillespie simulations for varying system size N (see [Fig 2A](#)). An excellent match between microscopic and macroscopic dynamics can be observed for every system size, for both the success rate \mathcal{S} and the convergence time \mathcal{C} . When




Fig 2. Collective decisions on a fully-connected network: comparison between the micro and the macro dynamics and scaling properties. (A,C) Comparison between the stochastic finite-size macroscopic model (black lines) and the multiagent implementation with both the homogeneous strategy (red lines) and the heterogenous strategy (green lines). Results are displayed for varying system size N . For each possible configuration (v_A, v_B) , 500 independent runs are performed. We show results for configurations with effectivity $\mathcal{E} > 0.7$. The plot is divided in two parts: in the bottom-right triangle, we consider the success rate S for each configuration (v_A, v_B) , where $v_A \geq v_B$. For each group size N , we show the isolines at $S = 0.9$. The gray triangle indicates quality value pairs below the target resolution $R = 0.15$ (i.e., configurations in which the two options are considered equivalent). In the top-left half of the plot, we consider the convergence time C for each configuration (v_A, v_B) , where $v_B \geq v_A$ (i.e., the symmetric problems with respect to the bottom-right plot). For each group size N , we show the isolines at $C = \hat{C}$. (B,D) Scaling of the convergence time C with the system size N . For each configuration (v_A, v_B) , we fit the curve $C = b N^a$ and we show the heat-map for the fitted coefficient a (see the bottom-right triangle showing the coefficient value for each configuration (v_A, v_B) , $v_A \geq v_B$) and b (see the top-left triangle showing the coefficient value for symmetric configurations (v_A, v_B) , $v_B \geq v_A$) across the decision space. Also in this case we show only configurations where $\mathcal{E} > 0.7$ for all N , and the white space indicates configurations with low effectivity. (A,B) Results for case-study 1A with $v_i \in [0, 1]$, $\gamma_i = 0.6v_i$, $\alpha_i = 0$, $\rho_i = 0.1v_i$, $\sigma_i = 1$ and $i \in \{A, B\}$, $\hat{C} = 50$ s. (C,D) Results for case-study 1B with $v_i \in [1, 10]$, $\gamma_i = \rho_i = v_i$, $\alpha_i = 1/v_i$, $\sigma_i = 10$ and $i \in \{A, B\}$, $\hat{C} = 1$ s.

doi:10.1371/journal.pone.0140950.g002

the difference in quality between the two options is above the resolution $R = 0.15$, the correct decision is taken in at least 90% of the cases (i.e., $S = 0.9$, as evidenced by the isolines in the bottom-right part of Fig 2A laying within the grey shaded area) for every system size but $N = 10$. Indeed, small groups suffer from stochastic fluctuations, reflected by a substantially lower success rate with respect to larger groups. Conversely, the speed of the process is lower for larger groups, as indicated by the isolines for $\hat{C} = 50$ s in the top-left part of Fig 2A. To quantify the scaling properties with respect to the system size, we analysed the convergence time for each decision problem as a function of N . We found a generalised adherence with a power law behaviour $C = b N^a$, with exponent $a \approx 0.2$ as shown in Fig 2B for macroscopic Gillespie simulations, and S2 Fig for multiagent simulations. The coefficient b also varies with the decision problem: the lower the option quality difference, the higher the coefficient. Looking at Fig 2B and S2 Fig, we observe that C scales similarly across different decision problems, with the exception of problems characterised by similar qualities (i.e., $v_A \approx v_B$) that require in

general more time for convergence. Finally, we show in [S3 Fig](#) an example of the convergence dynamics for a specific decision problem (e.g., $v_A = 0.9$, $v_B = 0.6$), which highlights the close correspondence between ODEs, Gillespie and multiagent simulations.

Case study I-B is concerned with the general case of value-sensitive decision making [23], and discusses the implementation in case of less restricting conditions with respect to the previous case study. We consider a quality range $v \in [1, 10]$, we fix the quorum for the collective decision to $\Psi_q = 0.8$ and we limit the total execution time to $T = 40$ s. Here, we also demonstrate a fully heterogeneous implementation of the multiagent system.

To obtain value-sensitive decision making [23], discovery and recruitment rates are assumed to be linearly proportional to the option quality v_i (i.e., $\gamma_i = \rho_i = v_i$, $i \in \{A, B\}$), the abandonment rate is inversely proportional to quality (i.e., $\alpha_i = 1/v_i$), while the cross-inhibition rate is constant ($\sigma_i = \bar{\sigma}$), which we fix to $\bar{\sigma} = 10$. Given the macroscopic parameterisation, we follow the same methodology described above to analyse the finite-size effects produced by the system size N , and we design the multiagent system following both the homogeneous and the heterogeneous strategies (see [Methods](#)). In the latter case, we use fixed thresholds for all transition probabilities and discuss the error introduced with respect to the macroscopic dynamics. According to the design pattern prescriptions, agents are updated every $\tau = 20$ ms (see [S1 Text](#)).

[Fig 2C](#) shows the match between the macroscopic Gillespie simulations and the multiagent implementation with both the homogeneous and the heterogeneous strategy, for varying system size N (see also [S3 Fig](#) for an example of convergence dynamics). The correspondence between macroscopic model and microscopic implementation is remarkable also in this case. Even with the rough approximation of fixed thresholds for the heterogeneous case, we note a good match of the micro-macro dynamics, although not perfect especially for the convergence time \mathcal{C} at low qualities. The results show that the studied parameterisation allows to reliably take decisions for above-resolution decision problems already with $N = 100$, as indicated by the success rate \mathcal{S} in the bottom-right part of [Fig 2C](#). Conversely, the convergence time \mathcal{C} is very similar across different system sizes. Also in this case, we analysed the scaling behaviour of the convergence time, and found adherence with a power law behaviour $\mathcal{C} = bN^a$, but with a very low exponent a (see [Fig 2D](#) for Gillespie simulations and [S2 Fig](#) for multiagent simulations). With the proposed parameterisation, \mathcal{C} becomes nearly independent of the system size N in large parts of the problem space. The coefficient b is rather low too, indicating fast decisions especially for large differences in quality. This is a result of the higher transition rates chosen for the macroscopic model, which are reflected by a smaller timestep τ in the multiagent implementation as prescribed by the design pattern. Finally, we studied the micro-macro link in a best-of- n scenario. The results presented in [S4 Fig](#) reveal a very good correspondence between multiagent and Gillespie simulations, therefore validating the methodology beyond the binary decision problems presented above.

Case study II: Collective decisions in a search and exploitation problem

Here, we present a case study that includes as main features spatiality and local interactions between agents. The experimental scenario is similar to the swarm robotics study of [16]. Point-size agents move in a 2D environment characterised by three areas—a *home* and two *target* areas—and must select the best quality target (see [Methods](#)). Uncommitted agents explore the environment to discover new options. Committed agents recruit and cross-inhibit other agents and periodically re-estimate the quality of the option they are committed to. Differently from [16], here quality is independent of distance. However, spatiality may influence the decision dynamics (e.g., the rate of discovery is higher for closer targets), and only an accurate design of the agent behaviour can lead to the systematic choice of the best available option.

We choose a macroscopic parameterisation similar to the one of case study I-B, but option quality v_i varies in $[0.1, 1]$ and also the cross-inhibition rate varies linearly with quality (i.e., $\sigma_i = v_i$). Besides depending on option quality, discovery is episodic (see [S1 Text](#)), being determined by the diffusive motion pattern of uncommitted agents that start searching from home. We model the macroscopic discovery rate to be proportional to the quality v_i and to decay with the target distance d_i as follows:

$$\gamma_i = \frac{v_i \mu e^{-\xi d_i}}{d_i} \tag{8}$$

where ξ and μ are parameters estimated from preliminary experiments ($\mu \approx 0.12$ and $\xi \approx 0.24$), although geometrical approximations could be used as well.

Besides discovery, spatiality influences also the interaction patterns among agents, given that interactions are possible only with agents in the local neighbourhood (see [Methods](#)). To ensure a well-mixed system and comply with the design pattern requirements, we limit interactions within the home area and we force agents to periodically return home. As described above, the design pattern prescribes to have fixed probabilities to become interactive (e.g., return home with probability \mathcal{P}_I) or latent (e.g., leave home with probability \mathcal{P}_L). Therefore, we designed the individual behaviour after the microscopic description of [Fig 1C](#) following both the homogeneous and the heterogeneous strategy, and we set $\mathcal{P}_I = 0.001$ and $\mathcal{P}_L = 9 \mathcal{P}_I$ to ensure a fraction $\eta_I = 0.1$ of interactive agents on average.

The design choices detailed above allow us to determine the microscopic parameterisation starting from the desired macroscopic transition rates (see [S5 Fig](#) for a comparison between the rates determined by design and those estimated from multiagent simulations through survival analysis [29]). We have tested the micro-macro link varying both option quality and target area distance, to observe how the implementation deals with the inclusion of spatial factors and interactive-latent dynamics. Indeed, the macroscopic model does not consider such factors, exception made for the model of discovery of [Eq \(8\)](#). In particular, we are interested in making consistent choices notwithstanding the target area distance. [Fig 3A](#) shows one such case for a homogeneous system in which the better option is also the farthest. The macroscopic model predicts convergence on the best-quality option (B in this case, see the trajectory starting from $\Psi_U = 1$), and the simulations are centred at the predicted stable point. Good agreement between macroscopic Gillespie simulations and multiagent implementation is observable also for the success rate \mathcal{S} (see [Fig 3A](#) inset). For same quality options, the target area distance biases the choice towards the closer target area (see [S6 Fig](#) for the homogeneous case and [S7 Fig](#) for the heterogeneous case). When both distance and option quality are equal (i.e., a completely symmetric condition), the system converges toward the one or the other option with equal probability, as shown in [Fig 3B](#) for the homogeneous case. Here too, the adherence between microscopic and macroscopic dynamics is remarkable. All these tests have been performed with $N = 500$ agents. Good agreement is observed also for different system sizes, as shown in [S8 Fig](#).

Discussion

The design pattern methodology we propose provides a complete framework that allows to move from the choice of the macroscopic parameterisation down to the implementation of the individual behaviour. Each step is supported by the principled understanding of the causal relationship between microscopic choices and macroscopic effects. We have substantiated the methodology with case studies that, despite being idealised, contain all the ingredients to be taken as reference for practical applications. In this respect, the inclusion of latent states for individual agents is particularly important, as it allows to preserve the micro-macro link also

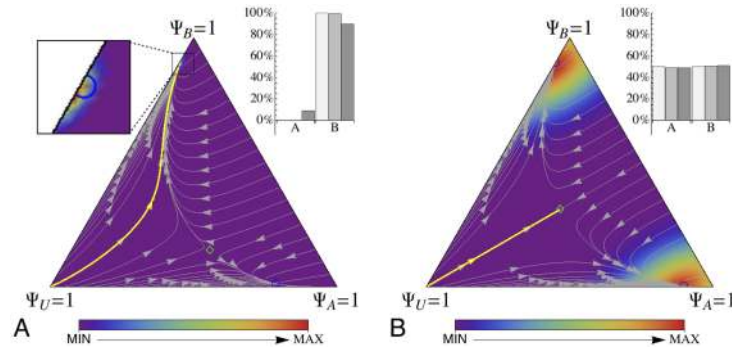


Fig 3. Collective decisions in a search and exploration problem: comparison between the micro and the macro dynamics. The state space of the system is presented as a ternary plot characterised by $\Psi_U + \Psi_A + \Psi_B = 1$, so that vertices correspond to fully-uncommitted or fully-committed populations. Macroscopic dynamics are indicated by trajectories and equilibrium points from the ODE model of Eq (1), parametrised according to the specific configuration. The bold yellow trajectory indicates the behaviour starting from a fully-uncommitted population ($\Psi_U = 1$). Stable equilibrium points are indicated as blue empty circles, while unstable points are indicated as green empty diamonds. The density map in the background represents the results of homogeneous multiagent simulations (1000 runs). The inset shows the success rate S for macroscopic Gillespie simulations (white bars) and multiagent simulations (homogeneous in light gray and heterogeneous in dark grey). (A) Micro-macro link for a decision problem in which the best option is also the farthest one ($v_A = 0.7 < v_B = 1$ and $d_A = 1.5 \text{ m} < d_B = 2.5 \text{ m}$). The magnify-glass effect allows to appreciate the close correspondence between the stable point predicted by the macroscopic model and the results from the multiagent simulations. (B) Micro-macro link for a completely symmetric decision problem ($v_A = v_B = 1$ and $d_A = d_B = 2.5 \text{ m}$).

doi:10.1371/journal.pone.0140950.g003

when interactions are sporadic or when spatiality interferes with the well-mixed assumption [16]. Indeed, departures from the macroscopic predictions are expected in case of heterogeneous interaction topologies, as is the case for other ordering processes [30, 31]. In this case, the micro-macro link could be preserved through the inclusion of heterogeneous mean-field approximations, which can correct the departure from the assumed well-mixed condition at the macroscopic level [32]. Future work should also take into account the macroscopic effects of interactions over adaptive and multi-layer networks, in order to (i) take into consideration the variability of the topology of interactions with time [33] and (ii) allow for the existence of different layers of connectivity among agents, each pertaining to specific context-related properties [34].

Besides engineering, our results can be relevant for better understanding the behaviour of natural systems. With respect to honeybee nest-site selection, our results provide testable hypotheses about the implementation algorithm employed by individual bees in relation to the proposed macroscopic model [22, 23]. Our implementation algorithm is simpler than other individual-based approaches [35, 36], as it abstracts several details that require assumptions difficult to be verified experimentally. Similarly to previous studies, [36, 37], we have shown in the second case study that spatiality affects the outcome of the collective decision biasing it towards closer sites. In our implementation, this is mainly the result of the quicker discovery of closer sites, while different latencies (e.g., shorter travel times) play a negligible role, thanks to the fixed probability of becoming interactive required to preserve the micro-macro link (see S1 Text). Field experiments should be targeted to verify the existence of a tradeoff between distance and quality. Furthermore, the effects of finite-sized groups in the decision dynamics predicted by the stochastic macroscopic simulations adhere with studies about group-size effects in natural conditions [38]: the larger the swarm the more accurate its decision. The scaling of decision time with group size that we have highlighted here represents another interesting aspect to investigate with field experiments. Finally, the extent to which behavioural

heterogeneities influence honeybee nest-site selection would be interesting to study, as genetic and molecular determinants of honeybee behaviour seem to play an important role [39, 40].

Behaviour heterogeneity in social systems is an important aspect not to be overlooked, as it can lead to interesting collective dynamics that are not attainable in fully homogeneous systems [41–43]. In our study, the choice of response thresholds for the heterogeneous implementation strategy is supported by the large literature on inter-individual variability in social insects [25, 44–47]. Recent studies have recognised the importance of including individual differences in behaviour—often referred to as personality or behavioural syndrome [48]—to better understand the collective dynamics [49, 50]. Here, we have highlighted the relationship between the distribution of individual thresholds and the collective response function, so that macroscopic predictions could be matched against estimates of the real threshold distribution [51]. We have also shown that fixed response thresholds well approximate the macroscopic dynamics especially for interactive processes like recruitment [25]. Associating fixed response thresholds with variable probability and intensity of responses may lead to more flexible and robust behaviour at the colony level [47]. Response thresholds well adhere with adaptive mechanisms for threshold adaptation [51, 52], allowing to finely tune the macroscopic response to match the statistical regularities that characterise the task. This adaptivity can result from evolutionary factors [53] as well as from development and learning [47]. Integrating adaptive mechanisms in the microscopic implementation could lead to improved performance [54, 55], and represents a natural extension for the proposed design pattern.

Methods

Multiagent simulations on fully-connected networks

We implemented a synchronous simulation for a multiagent system on a fully-connected network that directly derives from the design pattern implementation guidelines. At simulation start, each agent a_g estimates the quality \hat{v}_i of all available options $i \in \{1, \dots, n\}$, and on that basis computes its own transition probabilities $P_{\lambda, g}(\hat{v}_i)$, with $\lambda \in \{\gamma, \alpha, \rho, \sigma\}$. In the homogeneous case, these are computed in the same way for each agent according to the desired parameterisation, as prescribed by Eq (3). In the heterogeneous case, each agent a_g draws a random threshold δ_g from the random distribution D_λ and computes the transition probabilities as prescribed by Eq (4).

The simulation proceeds in discrete time steps of length τ . At each time step t , every agent updates its state following the PFSM of Fig 1A. All spontaneous transitions are always available. Conversely, interactive transitions depend on the interaction with a randomly selected partner, who shares its own commitment state and probabilities of recruitment and cross-inhibition. Given the well-mixed property ensured by the fully-connected topology, the population-dependent probabilities \mathcal{P}_{ψ_i} are estimated by randomly choosing a different agent $a_{\hat{g}}$ as partner at each time step and checking its state: transitions are activated if the selected partner is committed to some option i (see S1 Text for details). In this case, the probability of recruitment $P_{\rho, \hat{g}}$ and of cross-inhibition $P_{\sigma, \hat{g}}$ are received from the selected partner, otherwise they are null. In this way, the agent a_g has complete information to update its commitment state.

Particle-based multiagent simulations

Agents are point-size particles capable to move in a 2D environment. Movement is simulated through kinematic equations on the basis of the current agent speed v and orientation θ . The environment is an infinite plane and does not contain any wall or obstacle. No collision or physical interference among agents is taken into account, and agents are free to move in any

location of the 2D plane. The environment contains three circular areas with radius $r = 0.3$ m: home, target A and target B. Target areas are located at a variable distance $d_i \in [1.5 \text{ m}, 3.5 \text{ m}]$, $i \in \{A, B\}$, from the home area. Each target area is further characterised by a quality $v_i \in [0.1, 1]$; each agent can individually estimate the target area quality when is inside the area.

Agents move at a constant speed $v = 0.01$ m/s and communicate with their neighbours within an interaction range $d_I = 0.6$ m. Every $\tau = 1$ s, agents update both their state and their motion direction θ as detailed below.

Interactive and latent agents. As a consequence of local communications and of the distance between target areas, agents committed to different targets and uncommitted agents cannot always interact with each other. To ensure a well-mixed system, we limit interactions only when agents are within the home area. Agents remain in the home area for a time interval τ_I exponentially distributed with average $1/\mathcal{P}_I$. When a timeout expires, agents get latent and leave the home area. Agents remain latent for a time interval τ_L exponentially distributed with average $1/\mathcal{P}_L$. The motion pattern is conceived to ensure that agents are within the home area once they get interactive again (see below). To ensure that on average 10% of the agents are interactive within the home area, we set $\mathcal{P}_L = 9\mathcal{P}_I$.

Motion pattern. The agent motion direction θ is determined by the current agent state. While moving, odometry sensors are exploited to track the position of known areas. In this case study, we model noiseless sensors. In a more realistic implementation, agent-to-agent communication can be exploited to compensate for odometry noise [56].

An uncommitted agent a_g ($C(a_g) = C_U$) explores the environment in search of target areas: when latent, it chooses a random direction and leaves home moving in a straight line. Upon finding target area i , the agent makes an estimate of the quality \hat{v}_i and gets committed with probability $P_{\rho,g}(\hat{v}_i)$ (which is computed according to the homogeneous or heterogeneous implementation strategy). Then, it stores the target area estimated quality and location—which is kept updated through odometry—and returns back home. If an uncommitted agent has explored the environment for $[\tau_L/2]$ s without encountering any target area, it returns home to ensure that it gets interactive when already within the home area.

An agent a_g committed to option i ($C(a_g) = C_i$), while interactive, randomly moves within the home area to communicate with neighbours. When latent, it leaves home to return to the chosen target area i and re-estimate its quality \hat{v}_i . The agent returns home after $\tau_L/2$ s to become interactive within the home area. While latent, the committed agent abandons commitment with probability $P_{\alpha,g}(\hat{v}_i)$ and returns home. This probability is also computed according to either the homogeneous or the heterogeneous implementation strategy.

Interaction pattern. When interactive, all agents located in the home area can exchange short communication messages with a randomly chosen neighbour. These messages contain information on the agent state, which is used to estimate the population-dependent probabilities (see [S1 Text](#) for details).

When committed to option i , agent a_g also communicates the stored location of the target area and its own probability of recruitment $P_{\rho,g}(\hat{v}_i)$ and cross-inhibition $P_{\sigma,g}(\hat{v}_i)$, computed either with the homogeneous or the heterogeneous implementation strategy. Recruitment takes place if the partner is uncommitted, otherwise cross-inhibition takes place only if the distance between the target area locations internally stored by the interacting agents is larger than the target area radius r . In this way, cross-inhibition takes place only between agents committed to different options, as prescribed by the design pattern.

When the population is divided between interactive and latent agents, the design pattern prescribes that the dynamics of activity change must be faster than changes in the commitment state. To achieve this, we let agents interact only upon getting latent. By doing so, we guarantee

that the fraction of interactive agents committed to option i is always an unbiased representation of the global fraction Ψ_i (see [S1 Text](#)).

Supporting Information

S1 Text. Detailed description of the proposed design pattern methodology. In this document, we detail and formalise all aspects of the design pattern for decentralised decision making, in some cases repeating and extending the information succinctly given in the main text. (PDF)

S2 Text. Stability analysis and parameterisation choice for multiagent simulations on fully-connected networks. In this document, we provide additional results supporting case study I-A, in particular concerning the stability analysis of the macroscopic system and the selection of the desired macroscopic parameterisation. (PDF)

S1 Fig. Probabilistic Finite State Machines (PFSMs) describing the microscopic behaviour of an agent in average. Here, the notation $\mathcal{P}_i, \lambda \in \{\gamma, \alpha, \rho, \sigma\}, i \in \{1, \dots, n\}$ is a shorthand for $\mathcal{P}_i(v_i)$. (A) PFSM describing the basic commitment dynamics for n possible options. Spontaneous transitions are represented by solid lines, while interactive transitions are represented by dashed lines. (B) PFSM describing the coupled commitment and activity dynamics. Latent states are indicated in grey, and dash-dotted lines represent changes between latent and interactive states. (TIFF)

S2 Fig. Scaling of the convergence time \mathcal{C} with the system size N . For each configuration (v_A, v_B) , $v_A > v_B$ —and the symmetric case (v_A, v_B) , $v_B > v_A$ —we fit the curve $\mathcal{C} = bN^a$ and we show the heat-map for the fitted coefficient a (bottom-right) and b (top-left) across the decision space (see [Fig 2](#) for details). Also in this case we show only configurations where $\mathcal{E} > 0.7$. (A,B) Results for case study I-A with $v_i \in [0, 1]$, $\gamma_i = 0.6v_i$, $\alpha_i = 0$, $\rho_i = 0.1v_i$, $\sigma_i = 1$ and $i \in \{A, B\}$ for the homogenous (A) and the heterogeneous (B) implementation. (C,D) Results for case study I-B with $v_i \in [1, 10]$, $\gamma_i = \rho_i = v_i$, $\alpha_i = 1/v_i$, $\sigma_i = 10$ and $i \in \{A, B\}$ for the homogenous (C) and the heterogeneous (D) implementation. (TIFF)

S3 Fig. Mean trajectory of the population fractions Ψ_A, Ψ_B, Ψ_U over time from the initial condition $(\Psi_A = 0, \Psi_B = 0, \Psi_U = 1)$. Comparison at various levels of abstractions: mean field model (solid lines), macroscopic, finite-size Master equation (circles), and multiagent simulations (triangles). Simulations results are averaged over 100 independent runs. Errorbars are smaller than the symbols size, and are not displayed. (Top) Parameterisation of case study I-A with homogenous multiagent implementation and $v_A = 0.9$, $v_B = 0.6$. (Bottom) Parameterisation of case study I-B with heterogeneous multiagent implementation and $v_A = 9$, $v_B = 6$. (TIFF)

S4 Fig. Micro-macro link with varying number of options for case study I-B. We compare the macroscopic dynamics predicted by the mean-field model, the finite-size macroscopic dynamics simulated by the Gillespie algorithm and the microscopic dynamics resulting from homogeneous multiagent simulations ($N = 500$ agents). We fix the best option (A) to the maximum quality $v_A = 1$, and all other options to the same, lower quality v_i . The plot shows the fraction of the population committed to option A at the end of the simulation, plotted against the lower option quality v_i . Solid lines show the macroscopic prediction of the ODE system. The box-and-whiskers plots represent the statistics from Gillespie and multiagent simulations.

Boxes represent the inter-quartile range of the data (2000 runs), while the horizontal lines inside the boxes mark the median values. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range. Outliers are not shown. A very good match can be appreciated between microscopic and macroscopic dynamics, therefore validating the design pattern for best-of- n scenarios.

(TIFF)

S5 Fig. Comparison of the macroscopic transition rates resulting from design choices (solid lines) and estimated from the implemented multiagent system (points). The case for distance $d_A = d_B = 2.5$ m is shown. Estimates have been obtained through survival analysis computing the Nelson-Haalen estimator for the permanence time of agents in each state [29]. Survival analysis provides powerful non-parametric methods to estimate how the probability of events changes over time directly from the experimental data. See also [16] for details.

(TIFF)

S6 Fig. Comparison of the macroscopic dynamics with the microscopic homogeneous implementation for all tested decision problems (1000 runs for each setup). We vary both distance $d_A \in \{1.5, 2, 2.5, 3, 3.5\}$ m (from top to bottom) and quality $v_a \in \{0.1, 0.4, 0.7, 1.0\}$ (from left to right), while we keep fixed the distance $d_B = 2.5$ m and the quality $v_B = 1$.

(TIFF)

S7 Fig. Comparison of the macroscopic dynamics with the microscopic heterogeneous implementation for all tested decision problems (1000 runs for each setup). We vary both distance $d_A \in \{1.5, 2, 2.5, 3, 3.5\}$ m (from top to bottom) and quality $v_a \in \{0.1, 0.4, 0.7, 1.0\}$ (from left to right), while we keep fixed the distance $d_B = 2.5$ m and the quality $v_B = 1$.

(TIFF)

S8 Fig. Micro-macro link with varying group size for the search and exploitation task. We show results for $v_A = 0.7$, $v_B = 1.0$, $d_B = 2.5$ m and $d_A \in [1.5, 3.5]$ m. In this condition, B is the option to be selected, therefore we show the resulting fraction Ψ_B as the system size N varies in [50, 500]. The plot shows results from 2000 simulations for each configuration. Dashed lines represent the equilibrium point of the mean-field model, white boxes represent the master equation approximated through Gillespie simulations, light grey boxes represent homogeneous multiagent simulations, dark grey boxes represent heterogeneous multiagent simulations. Boxes represent the inter-quartile range of the data (2000 runs), while the horizontal lines inside the boxes mark the median values. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range. Empty circles mark the outliers. A good match between macroscopic and microscopic implementation can be appreciated for every group size. The larger discrepancies are recognised with the heterogeneous implementation for the most difficult decision problem (e.g., the wrong option is also the closest one: $d_A = 1.5$).

(TIFF)

Author Contributions

Conceived and designed the experiments: AR VT. Performed the experiments: AR. Analyzed the data: AR VT. Contributed reagents/materials/analysis tools: AR GV CFO MD VT. Wrote the paper: AR GV CFO MD VT.

References

1. Liu YY, Slotine JJ, Barabási AL. Controllability of complex networks. *Nature*. 2011; 473(7346):167–173. PMID: [21562557](https://pubmed.ncbi.nlm.nih.gov/21562557/)

2. Helbing D. Globally networked risks and how to respond. *Nature*. 2014; 497(7447):51–59. doi: [10.1038/nature12047](https://doi.org/10.1038/nature12047)
3. Lee EA, Hartmann B, Kubiawicz J, Simunic Rosing T, Wawrzynek J, Wessel D, et al. The Swarm at the Edge of the Cloud. *IEEE Design & Test*. 2014; 31(3):8–20. doi: [10.1109/MDAT.2014.2314600](https://doi.org/10.1109/MDAT.2014.2314600)
4. Couzin ID. Collective cognition in animal groups. *Trends Cogn Sci*. 2009; 13(1):36–43. doi: [10.1016/j.tics.2008.10.002](https://doi.org/10.1016/j.tics.2008.10.002) PMID: [19058992](https://pubmed.ncbi.nlm.nih.gov/19058992/)
5. Baronchelli A, Ferrer-i-Cancho R, Pastor-Satorras R, Chater N, Christiansen MH. Networks in Cognitive Science. *Trends Cogn Sci*. 2013; 17(7):348–360. doi: [10.1016/j.tics.2013.04.010](https://doi.org/10.1016/j.tics.2013.04.010) PMID: [23726319](https://pubmed.ncbi.nlm.nih.gov/23726319/)
6. Brambilla M, Ferrante E, Birattari M, Dorigo M. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*. 2013; 7(1):1–41. doi: [10.1007/s11721-012-0075-2](https://doi.org/10.1007/s11721-012-0075-2)
7. Akyldiz IF, Lo BF, Balakrishnan R. Cooperative spectrum sensing in cognitive radio networks: A survey. *Physical Communication*. 2011; 4(1):40–62. doi: [10.1016/j.phycom.2010.12.003](https://doi.org/10.1016/j.phycom.2010.12.003)
8. Derler P, Lee EA, Sangiovanni Vincentelli A. Modeling Cyber-Physical Systems. *Proceedings of the IEEE*. 2012; 100(1):13–28. doi: [10.1109/JPROC.2011.2160929](https://doi.org/10.1109/JPROC.2011.2160929)
9. Hamann H, Wörn H. A framework of space–time continuous models for algorithm design in swarm robotics. *Swarm Intelligence*. 2008; 2(2):209–239. doi: [10.1007/s11721-008-0015-3](https://doi.org/10.1007/s11721-008-0015-3)
10. Michael N, Kumar V. Planning and Control of Ensembles of Robots with Non-holonomic Constraints. *Int J Rob Res*. 2009; 28(8):962–975. doi: [10.1177/0278364909340280](https://doi.org/10.1177/0278364909340280)
11. Berman S, Kumar V, Nagpal R. Design of control policies for spatially inhomogeneous robot swarms with application to commercial pollination. In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE; 2011. p. 378–385.
12. Sartoretti G, Hongler MO, de Oliveira ME, Mondada F. Decentralized self-selection of swarm trajectories: from dynamical systems theory to robotic implementation. *Swarm Intelligence*. 2014; 8(4):329–351. doi: [10.1007/s11721-014-0101-7](https://doi.org/10.1007/s11721-014-0101-7)
13. Gamma E, Helm R, Johnson R, Vlissides J. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley Professional, Boston, MA; 1995.
14. Babaoğlu G, Canright G, Deutsch A, Di Caro GA, Ducatelle F, Gambardella LM, et al. *Design Patterns from Biology for Distributed Computing*. *ACM Transactions on Autonomous and Adaptive Systems*. 2006; 1(1):26–66. doi: [10.1145/1152934.1152937](https://doi.org/10.1145/1152934.1152937)
15. Fernandez-Marquez JL, Di Marzo Serugendo G, Montagna S, Viroli M, Arcos JL. Description and composition of bio-inspired design patterns: a complete overview. *Natural Computing*. 2013; 12(1):43–67. doi: [10.1007/s11047-012-9324-y](https://doi.org/10.1007/s11047-012-9324-y)
16. Reina A, Miletitch R, Dorigo M, Trianni V. A quantitative micro-macro link for collective decisions: The shortest path discovery/selection example. *Swarm Intelligence*. 2015; 9(2–3):75–102. doi: [10.1007/s11721-015-0105-y](https://doi.org/10.1007/s11721-015-0105-y)
17. Halloy J, Sempo G, Caprari G, Rivault C, Asadpour M, Tache F, et al. Social integration of robots into groups of cockroaches to control self-organized choices. *Science*. 2007; 318(5853):1155. doi: [10.1126/science.1144259](https://doi.org/10.1126/science.1144259) PMID: [18006751](https://pubmed.ncbi.nlm.nih.gov/18006751/)
18. Vigeliu M, Meyer B, Pascoe G. Multiscale Modelling and Analysis of Collective Decision Making in Swarm Robotics. *PLoS ONE*. 2014; 9(11):e111542–19. doi: [10.1371/journal.pone.0111542](https://doi.org/10.1371/journal.pone.0111542) PMID: [25369026](https://pubmed.ncbi.nlm.nih.gov/25369026/)
19. Srivastava V, Leonard NE. Collective Decision-Making in Ideal Networks: The Speed-Accuracy Trade-off. *IEEE Transactions on Control of Network Systems*. 2014; 1(1):121–132. doi: [10.1109/TCNS.2014.2310271](https://doi.org/10.1109/TCNS.2014.2310271)
20. Valentini G, Hamann H, Dorigo M. Efficient Decision-Making in a Self-Organizing Robot Swarm: On the Speed Versus Accuracy Trade-Off. In: *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*. International Foundation for Autonomous Agents and Multiagent Systems; 2015. p. 1305–1314.
21. Marshall JAR, Bogacz R, Dornhaus A, Planqué R, Kovacs T, Franks NR. On optimal decision-making in brains and social insect colonies. *J R Soc Interface*. 2009; 6(40):1065–74. doi: [10.1098/rsif.2008.0511](https://doi.org/10.1098/rsif.2008.0511) PMID: [19324679](https://pubmed.ncbi.nlm.nih.gov/19324679/)
22. Seeley TD, Visscher PK, Schlegel T, Hogan PM, Franks NR, Marshall JAR. Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*. 2012; 335(6064):108–11. doi: [10.1126/science.1210361](https://doi.org/10.1126/science.1210361) PMID: [22157081](https://pubmed.ncbi.nlm.nih.gov/22157081/)
23. Pais D, Hogan PM, Schlegel T, Franks NR, Leonard NE, Marshall JAR. A Mechanism for Value-Sensitive Decision-Making. *PLoS ONE*. 2013; 8(9):e73216. doi: [10.1371/journal.pone.0073216](https://doi.org/10.1371/journal.pone.0073216) PMID: [24023835](https://pubmed.ncbi.nlm.nih.gov/24023835/)

24. Crespi V, Galstyan A, Lerman K. Top-down vs bottom-up methodologies in multi-agent system design. *Auton Robots*. 2008; 24(3):303–313. doi: [10.1007/s10514-007-9080-5](https://doi.org/10.1007/s10514-007-9080-5)
25. Robinson EJH, Franks NR, Ellis S, Okuda S, Marshall JAR. A Simple Threshold Rule Is Sufficient to Explain Sophisticated Collective Decision-Making. *PLoS ONE*. 2011; 6(5):e19981. doi: [10.1371/journal.pone.0019981](https://doi.org/10.1371/journal.pone.0019981) PMID: [21629645](https://pubmed.ncbi.nlm.nih.gov/21629645/)
26. Chittka L, Skorupski P, Raine NE. Speed-accuracy tradeoffs in animal decision making. *Trends Ecol Evol*. 2009; 24(7):400–407. doi: [10.1016/j.tree.2009.02.010](https://doi.org/10.1016/j.tree.2009.02.010) PMID: [19409649](https://pubmed.ncbi.nlm.nih.gov/19409649/)
27. Franks NR, Richardson TO, Stroeymeyt N, Kirby RW, Amos WMD, Hogan PM, et al. Speed-cohesion trade-offs in collective decision making in ants and the concept of precision in animal behaviour. *Animal Behaviour*. 2013; 85(6):1233–1244. doi: [10.1016/j.anbehav.2013.03.010](https://doi.org/10.1016/j.anbehav.2013.03.010)
28. Gillespie DT. Exact stochastic simulation of coupled chemical reactions. *J Phys Chem A*. 1977; 81(25):2340–2361. doi: [10.1021/j100540a008](https://doi.org/10.1021/j100540a008)
29. Nelson W. Hazard plotting for incomplete failure data. *Journal of Quality Technology*. 1969; 1:27–52.
30. Sood V, Redner S. Voter Model on Heterogeneous Graphs. *Phys Rev Lett*. 2005; 94(17):178701–4. doi: [10.1103/PhysRevLett.94.178701](https://doi.org/10.1103/PhysRevLett.94.178701) PMID: [15904343](https://pubmed.ncbi.nlm.nih.gov/15904343/)
31. Baronchelli A, Dall'Asta L, Barrat A, Loreto V. Topology-induced coarsening in language games. *Phys Rev E Stat Nonlin Soft Matter Phys*. 2006; 73(1):015102. doi: [10.1103/PhysRevE.73.015102](https://doi.org/10.1103/PhysRevE.73.015102) PMID: [16486202](https://pubmed.ncbi.nlm.nih.gov/16486202/)
32. Moretti P, Liu SY, Baronchelli A, Pastor-Satorras R. Heterogenous mean-field analysis of a generalized voter-like model on networks. *Eur Phys J B*. 2012; 85(3):1–6. doi: [10.1140/epjb/e2012-20501-1](https://doi.org/10.1140/epjb/e2012-20501-1)
33. Groß T, Blasius B. Adaptive coevolutionary networks: a review. *Journal of The Royal Society Interface*. 2008; 5(20):259–271. doi: [10.1098/rsif.2007.1229](https://doi.org/10.1098/rsif.2007.1229)
34. Boccaletti S, Bianconi G, Criado R, del Genio CI, Gómez-Gardeñes J, Romance M, et al. The structure and dynamics of multilayer networks. *Physics Reports*. 2014; 544(1):1–122. doi: [10.1016/j.physrep.2014.07.001](https://doi.org/10.1016/j.physrep.2014.07.001)
35. Passino KM, Seeley TD. Modeling and analysis of nest-site selection by honeybee swarms: the speed and accuracy trade-off. *Behav Ecol Sociobiol*. 2006; 59(3):427–442. doi: [10.1007/s00265-005-0067-y](https://doi.org/10.1007/s00265-005-0067-y)
36. Janson S, Middendorf M, Beekman M. Searching for a new home—scouting behavior of honeybee swarms. *Behav Ecol*. 2007; 18(2):384–392. doi: [10.1093/beheco/arl095](https://doi.org/10.1093/beheco/arl095)
37. Laomettachit T, Termsaithong T, Sae-Tang A, Duangphakdee O. Decision-making in honeybee swarms based on quality and distance information of candidate nest sites. *J Theor Biol*. 2015; 364:21–30. doi: [10.1016/j.jtbi.2014.09.005](https://doi.org/10.1016/j.jtbi.2014.09.005) PMID: [25218431](https://pubmed.ncbi.nlm.nih.gov/25218431/)
38. Schaerf TM, Makinson JC, Myerscough MR, Beekman M. Do small swarms have an advantage when house hunting? The effect of swarm size on nest-site selection by *Apis mellifera*. *J R Soc Interface*. 2013; 10(87):20130533–20130533. doi: [10.1098/rsif.2013.0533](https://doi.org/10.1098/rsif.2013.0533) PMID: [23904590](https://pubmed.ncbi.nlm.nih.gov/23904590/)
39. Mattila HR, Seeley TD. Genetic diversity in honey bee colonies enhances productivity and fitness. *Science*. 2007; 317(5836):362–364. doi: [10.1126/science.1143046](https://doi.org/10.1126/science.1143046) PMID: [17641199](https://pubmed.ncbi.nlm.nih.gov/17641199/)
40. Liang ZS, Nguyen T, Mattila HR, Rodriguez-Zas SL, Seeley TD, Robinson GE. Molecular Determinants of Scouting Behavior in Honey Bees. *Science*. 2012; 335(6073):1225–1228. doi: [10.1126/science.1213962](https://doi.org/10.1126/science.1213962) PMID: [22403390](https://pubmed.ncbi.nlm.nih.gov/22403390/)
41. Martino AD, Marsili M. Statistical mechanics of socio-economic systems with heterogeneous agents. *Journal of Physics A: Mathematical and General*. 2006; 39(43):R465–R540. doi: [10.1088/0305-4470/39/43/R01](https://doi.org/10.1088/0305-4470/39/43/R01)
42. Helbing D, editor. *Social Self-Organization. Understanding Complex Systems*. Springer Verlag, Berlin, Germany; 2012.
43. Huang K, Wang T, Cheng Y, Zheng X. Effect of Heterogeneous Investments on the Evolution of Cooperation in Spatial Public Goods Game. *PLoS ONE*. 2015; 10(3):e0120317–10. doi: [10.1371/journal.pone.0120317](https://doi.org/10.1371/journal.pone.0120317) PMID: [25781345](https://pubmed.ncbi.nlm.nih.gov/25781345/)
44. Burns JG. Impulsive bees forage better: the advantage of quick, sometimes inaccurate foraging decisions. *Anim Behav*. 2005; 70(6):e1–e5. doi: [10.1016/j.anbehav.2005.06.002](https://doi.org/10.1016/j.anbehav.2005.06.002)
45. Mailléux AC, Detrain C, Deneubourg JL. Starvation drives a threshold triggering communication. *J Exp Biol*. 2006; 209(21):4224–4229. doi: [10.1242/jeb.02461](https://doi.org/10.1242/jeb.02461) PMID: [17050837](https://pubmed.ncbi.nlm.nih.gov/17050837/)
46. Dussutour A, Nicolis SC, Despland E, Simpson SJ. Individual differences influence collective behaviour in social caterpillars. *Anim Behav*. 2008; 76(1):5–16. doi: [10.1016/j.anbehav.2007.12.009](https://doi.org/10.1016/j.anbehav.2007.12.009)
47. Jeanson R, Weidenmüller A. Interindividual variability in social insects—proximate causes and ultimate consequences. *Biol Rev Camb Philos Soc*. 2013 Dec; 89(3):671–687. doi: [10.1111/brv.12074](https://doi.org/10.1111/brv.12074) PMID: [24341677](https://pubmed.ncbi.nlm.nih.gov/24341677/)

48. Jandt JM, Bengston S, Pinter-Wollman N, Pruitt JN, Raine NE, Dornhaus A, et al. Behavioural syndromes and social insects: personality at multiple levels. *Biol Rev Camb Philos Soc.* 2014; 89(1):48–67. doi: [10.1111/brv.12042](https://doi.org/10.1111/brv.12042) PMID: [23672739](https://pubmed.ncbi.nlm.nih.gov/23672739/)
49. Wray MK, Seeley TD. Consistent personality differences in house-hunting behavior but not decision speed in swarms of honey bees (*Apis mellifera*). *Behav Ecol Sociobiol.* 2011; 65(11):2061–2070. doi: [10.1007/s00265-011-1215-1](https://doi.org/10.1007/s00265-011-1215-1)
50. Planas-Sitjà I, Deneubourg JL, Gibon C, Sempo G. Group personality during collective decision-making: a multi-level approach. *Proc Biol Sci.* 2015; 282(1802):20142515. doi: [10.1098/rspb.2014.2515](https://doi.org/10.1098/rspb.2014.2515) PMID: [25652834](https://pubmed.ncbi.nlm.nih.gov/25652834/)
51. Weidenmüller A. The control of nest climate in bumblebee (*Bombus terrestris*) colonies: interindividual variability and self reinforcement in fanning response. *Behav Ecol.* 2004; 15(1):120–128. doi: [10.1093/beheco/arg101](https://doi.org/10.1093/beheco/arg101)
52. Theraulaz G, Bonabeau E, Deneubourg JL. Response threshold reinforcements and division of labour in insect societies. *Proc Biol Sci.* 1998; 265(1393):327–332. doi: [10.1098/rspb.1998.0299](https://doi.org/10.1098/rspb.1998.0299)
53. Duarte A, Pen I, Keller L, Weissing FJ. Evolution of self-organized division of labor in a response threshold model. *Behav Ecol Sociobiol.* 2012; 66(6):947–957. doi: [10.1007/s00265-012-1343-2](https://doi.org/10.1007/s00265-012-1343-2) PMID: [22661824](https://pubmed.ncbi.nlm.nih.gov/22661824/)
54. Franklin EL, Robinson EJJ, Marshall JAR, Sendova-Franks AB, Franks NR. Do ants need to be old and experienced to teach? *J Exp Biol.* 2012; 215(8):1287–1292. doi: [10.1242/jeb.064618](https://doi.org/10.1242/jeb.064618) PMID: [22442366](https://pubmed.ncbi.nlm.nih.gov/22442366/)
55. Westhus C, Kleineidam CJ, Roces F, Weidenmueller A. Behavioural plasticity in the fanning response of bumblebee workers: impact of experience and rate of temperature change. *Anim Behav.* 2013; 85(1):27–34. doi: [10.1016/j.anbehav.2012.10.003](https://doi.org/10.1016/j.anbehav.2012.10.003)
56. Gutiérrez A, Campo A, Monasterio-Huelin F, Magdalena L, Dorigo M. Collective decision-making based on social odometry. *Neural computing & applications.* 2010; 19(6):807–823. doi: [10.1007/s00521-010-0380-x](https://doi.org/10.1007/s00521-010-0380-x)