

By force of habit: Self-trapping in a dynamical utility landscape

Cite as: Chaos **30**, 053123 (2020); <https://doi.org/10.1063/5.0009518>

Submitted: 31 March 2020 . Accepted: 25 April 2020 . Published Online: 11 May 2020

 José Moran, Antoine Fosset, Davide Luzzati, Jean-Philippe Bouchaud, and  Michael Benzaquen

COLLECTIONS

Paper published as part of the special topic on [Dynamics of Social Systems](#)



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

[Transport and concentration of wealth: Modeling an amenities-based-theory](#)

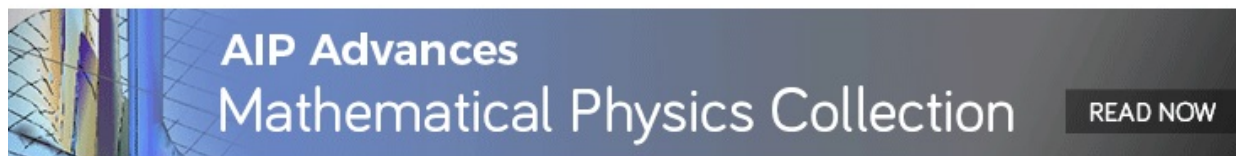
Chaos: An Interdisciplinary Journal of Nonlinear Science **30**, 053110 (2020); <https://doi.org/10.1063/5.0003767>

[Fear induced explosive transitions in the dynamics of corruption](#)

Chaos: An Interdisciplinary Journal of Nonlinear Science **30**, 063107 (2020); <https://doi.org/10.1063/5.0004826>

[Fat tails and black swans: Exact results for multiplicative processes with resets](#)

Chaos: An Interdisciplinary Journal of Nonlinear Science **30**, 033104 (2020); <https://doi.org/10.1063/1.5141837>



AIP Advances
Mathematical Physics Collection

READ NOW



By force of habit: Self-trapping in a dynamical utility landscape

Cite as: Chaos 30, 053123 (2020); doi: 10.1063/5.0009518

Submitted: 31 March 2020 · Accepted: 25 April 2020 ·

Published Online: 11 May 2020



View Online



Export Citation



CrossMark

José Moran,^{1,2,a)}  Antoine Fosset,^{2,3} Davide Luzzati,^{2,3} Jean-Philippe Bouchaud,^{2,4} and Michael Benzaquen^{2,3,4} 

AFFILIATIONS

¹Centre d'Analyse et de Mathématiques Sociales, EHESS, 54 Boulevard Raspail, 75006 Paris, France

²Chair of Econophysics and Complex Systems, Ecole polytechnique, 91128 Palaiseau Cedex, France

³LadHyX UMR CNRS 7646, Ecole polytechnique, 91128 Palaiseau Cedex, France

⁴Capital Fund Management, 23 Rue de l'Université, 75007 Paris, France

Note: This article is part of the Focus Issue, Dynamics of Social Systems.

^{a)}**Author to whom correspondence should be addressed:** jose.moran@polytechnique.org

ABSTRACT

Historically, rational choice theory has focused on the utility maximization principle to describe how individuals make choices. In reality, there is a computational cost related to exploring the universe of available choices and it is often not clear whether we are truly maximizing an underlying utility function. In particular, memory effects and habit formation may dominate over utility maximization. We propose a stylized model with a history-dependent utility function, where the utility associated to each choice is increased when that choice has been made in the past, with a certain decaying memory kernel. We show that self-reinforcing effects can cause the agent to get stuck with a choice by sheer force of habit. We discuss the special nature of the transition between free exploration of the space of choice and self-trapping. We find, in particular, that the trapping time distribution is precisely a Zipf law at the transition, and that the self-trapped phase exhibits *super-aging* behavior.

Published under license by AIP Publishing. <https://doi.org/10.1063/5.0009518>

In modeling the behavior of agents, the common view in standard economics is that their actions are guided by the maximization of an utility function. For convenience, the utility for each agent is often thought of as independent of the actions of others as well as static in time. A myriad of results in economic theory actually rest upon this assumption. Complexity economics has recently begun to tackle the issue of interactions between agents with analytical and numerical tools, and we address here the possibility of reinforcement mechanisms that make an agent's utility depend on his past. This simple toy model leads to non-ergodic dynamics, where the agent's actions depend crucially on past decisions.

I. INTRODUCTION

A key assumption in rational choice theory is that individuals set their preferences according to an utility maximization principle. Each choice an individual can make is assigned a certain “utility,” i.e., a quantity measuring the satisfaction it provides to the agent and frequently related to the dispassionate forecast of a related payoff.

This framework is often accompanied by the assumption that the agent considers *all* available choices present to her/him, weighs their utilities against one another, and then makes her/his choice taking into account possible constraints, such as a finite budget.

A number of criticisms to this view of human behavior have emerged, with, e.g., Simon¹ as a key figure highlighting that individuals may be “satisfiers” rather than pure optimizers, in the sense that there is both a computational cost and a cognitive bias related to considering the universe of available choices. Sometimes, finding the optimum of the utility function can itself be such a computationally hard problem that even the most powerful computers would not be able to find it in a reasonable amount of time. This led to the idea of bounded rationality as a way to model real agents.^{2–5} More recently, Kahneman^{6,7} pointed at what he considers to be significant divergences between economics and psychology in their assumptions of human behavior, with a special emphasis on the empirical evidence of the cognitive biases, and, therefore, the irrationality that guides individual behavior. A pervasive effect, for example, is that the utility of a certain choice strongly depends on the choice made by others. These so-called “externalities” can lead to interesting collective

effects, where choices made by agents synchronize and condense on a small subset of choices, or lead to confidence crises—see, for example, Refs. 8–12.

An interesting idea developed in Ref. 13 is the fact that the utility associated to a certain decision may depend also on our memory if it has already been made in the past. Here, we propose a simple model that encapsulates this idea and show that this too can lead to choices that do not necessarily conform to their “objective” utilities, but are rather dominated by past choices alone. This is related to what economists call “habit formation.”^{14–19} Memory effects chisel the utility landscape in a way that may render objectively sub-optimal choices subjectively optimal. In the case of sufficiently long-range memory, agents may, in a self-fulfilling kind of way, become “trapped” forever in a certain choice and stop exploring alternative choices.

As a practical example, one may imagine a situation where one must choose where to have lunch every day. Standard rational theory dictates that we ought to scrutinize every restaurant, eatery, and cafeteria, taking into account our personal tastes and the costs associated with going to any of these places. In contrast, we want to model the fact that habit may take over: instead of seeking to maximize a certain objective cost function, we are likely to persist in going to a specific place just because we are used to it. Anecdotal evidence shows that this is indeed what often happens in practice!

Our model assumes that the utility landscape is affected by past choices, with a memory kernel that decays with time. Agents can change their decision using a logit (or Metropolis) rule, parameterized by an “intensity of choice” β that plays the role of the inverse temperature in statistical physics. This type of model belongs to a wide class of so-called “reinforcement” models, which contains Polya Urns, Reinforced Random Walks, Elephant Walks, etc.—for a review, see Ref. 20, and references therein (for recent developments, see also Refs. 21 and 22). Such models have also gained traction in the economics literature, where positive reinforcement of certain choices made by agents are shown to impact the emergence of certain macro-outcomes and structures.^{23,24}

After properly defining our model, we provide analytical arguments to confirm the intuition that sufficiently strong memory effects, coupled with the optimization of the subjective memory-induced utility, can lead to “self-trapping,” i.e., the agent sticks to a choice whose objective utility is not necessarily maximal, simply by force of habit. We confirm our result via numerical simulations that explore different topologies for the space of different choices. We discover a particularly interesting dynamical transition when the memory kernel decays as the inverse of time, with rather unusual scaling and super-aging properties. We finally propose possible extensions.

II. A SIMPLE MODEL

Consider a set of N discrete choices, labeled $(x_i)_{1 \leq i \leq N}$, to which we assign an utility—a measure of the value an individual assigns a given choice. The perceived utility of site x_i and time t is postulated to be

$$U(x_i, t) = U_0(x_i) \left(1 + \sum_{t'=0}^t \phi(t-t') \mathbf{1}_{x(t')=x_i} \right), \quad (2.1)$$

where the first term on the right-hand side is the intrinsic, or objective utility of the choice, while the second accounts for memory effects, affecting the utility of that choice for the only reason that the individual has picked it in the past. [One may also think, in the physicist’s language, of an energy landscape (akin to minus the utility), where the energy of a given site or configuration increases or decreases if the system has already visited that site.] The decaying memory kernel ϕ encodes that more recent choices have a stronger effect, and $x(t)$ denotes the choice of the individual at time t . Hence, past history “chisels” the utility landscape, in a way similar to ants leaving a pheromone trace that guide other ants along the same path, or rivers creating their own bed through erosion. Note that in most reinforcement random walk models reviewed in Ref. 20, infinite memory span is assumed, i.e., $\phi(t) = \text{constant}$, while we will be mostly concerned here with decaying memory kernels.

The sign of the kernel ϕ separates two different cases: $\phi < 0$ indicates a situation where an individual grows weary of his past choices, while $\phi > 0$ corresponds to the case where an individual becomes increasingly endeared with them. In agreement with intuition, the former case leads to an exploration of all the choices unless the optimal choice has an utility too far apart from the rest to be sufficiently affected by the kernel. In all that follows, we focus on the more interesting case $\phi \geq 0$. The reason behind studying such utility reinforcement lies in the behavioral idea that people tend to prefer what they already know, thus paving the way for “habit formation” as in Ref. 25 (see also Refs. 11 and 15–19). We then consider the following dynamics. An individual, standing by choice x_i at time t , draws an alternative x_j from a certain ensemble of “nearby choices” ∂_i , e.g., the set of neighbors of i in a graph \mathcal{G} , with probability,

$$\mathcal{T}_{x_i \rightarrow x_j} = \frac{\mathbf{1}_{x_j \in \partial_i}}{\mathcal{N}_i}, \quad \mathcal{N}_i := \sum_j \mathbf{1}_{x_j \in \partial_i}, \quad (2.2)$$

where \mathcal{N}_i is the number of neighbors of i . Restricting to nearby choices is a parsimonious way to model adaptation costs that penalize large decision changes. However, our framework is quite versatile since the topology of the graph \mathcal{G} is arbitrary, and we will consider different cases below.

The target choice x_j is then adopted with the logit probability, standard in choice theory²⁶ (for non-trivial trapping to emerge, we consider graphs without singletons, that is to say that all sites have a non-empty set of neighbors that are different to itself),

$$p(x_i \rightarrow x_j) = \frac{1}{1 + e^{\beta[U(x_i, t) - U(x_j, t)]}}, \quad (2.3)$$

where β is called the “intensity of choice” and accounts for the degree of rationality (it is the analog of the inverse temperature in statistical mechanics). Indeed, as long as $0 < \beta < \infty$, the agent is more likely to switch whenever $U(x_j, t) > U(x_i, t)$ (optimizing behavior), but the probability to pick a choice with a lower utility is non-zero, which encodes for bounded-rationality (or uncertainty about the true utility) in the economics literature. In the $\beta \rightarrow 0$ limit (equivalent to the infinite temperature limit in physics), the agent explores the whole space of possible choices without taking their utility into account. In the opposite limit $\beta \rightarrow \infty$ (or zero temperature), the agent has a greedy behavior and only switches to choices with a higher utility, but this also implies that he/she may stay in a




FIG. 1. Schematic representation of our problem at a given time t . The plot on the top depicts the case of random “objective” utilities $U_0(x)$, while the one on the bottom shows the situation where they are uniform $U_0(x) = U_0$. In both plots, the solid black ball represents the choice made at time t , while the empty ball represents the choice made at time $t - 1$. Both correspond to a simulation run with a power-law kernel $\phi(t) \propto (1 + t)^{-\gamma}$ with $\gamma = 1.5$ and $\beta = 0.2$, on a fully connected graph.

local maximum instead of taking the chance to explore all available possibilities. An illustration of these dynamics is given in Fig. 1.

When $\phi = 0$, the dynamics is that of the Metropolis–Hastings algorithm used to sample the Boltzmann–Gibbs distribution.^{27,28} The stationary state of the dynamics is such that the probability to pick choice x_i is proportional to $N_i e^{\beta U_0(x_i)}$. This can by itself lead to interesting phenomena depending on the statistics of U_0 . For example, the study of the random energy model,^{29,30} for a finite value of N and for Gaussian utilities of variance σ^2 , shows that for $\beta > \beta_c = \sqrt{2 \ln N}/\sigma$, the probability measure condenses on a small number of choices, much smaller than N .

III. NON-ERGODICITY AND CONDENSATION OF CHOICES

Adding the kernel introduces the possibility that the agent gets stuck in a non-optimal choice exclusively through memory effects: staying a long time in a given choice self-reinforces its utility, thereby increasing the likelihood to stay there and leading to non-ergodic dynamics. To study the possible condensation or trapping induced by memory alone, we restrict ourselves to the case where $U_0(x_i) = 1, \forall i$ henceforth (any other value of U_0 can be reabsorbed in β). The interplay between memory-induced trapping and utility heterogeneity is quite interesting in itself, but we leave it to future investigations.

We consider an agent starting from a given choice x_0 at time $t = 0$ and follow his/her evolution for times $t = 1, \dots, T$ with T sufficiently large. We then compute the empirical state histogram $p_i = \sum_t \mathbf{1}_{x(t)=x_i}/T$ and define the order parameter h , in a similar way to the inverse participation ratio used in condensed matter physics^{31,32} or to the Herfindahl index in economics,³³ as

$$h := \sum_{i=1}^N p_i^2. \tag{3.1}$$

This parameter indicates how the agent has explored the space of possible choices: if all choices were visited with equal probability, then one has immediately $p_i = 1/N$ and thus $h = \sum_i 1/N^2 = 1/N$. On the other hand, if the agent was stuck in a single choice j , then $p_i = \delta_{i,j}$ and so $h = 1$. Therefore, $1/h$ gives an order of magnitude of the number of different choices picked by the agent during time T . In practice, we average h over a large number of trajectories and starting sites x_0 , to obtain an average parameter $\langle h \rangle$. For a set of simulations on a graph \mathcal{G} with N choices and lasting a time T , we, therefore, define the critical value β_c , defining the crossover between $h = \mathcal{O}(1/N)$ and $h = \mathcal{O}(1)$ as the value for β that maximizes the variance of h over different trajectories.

An important question is whether β_c corresponds to a true transition or to a mere crossover. This depends on the L^1 norm of the memory kernel, $|\phi| = \sum_{t=0}^{\infty} \phi(t)$. Suppose that this norm is finite. Then, if the agent has been stuck in a given site i for a time $t \gg 1$, we can approximate its utility by $(1 + |\phi|)$. The difference in utility with the neighboring choices thus remains finite. For any finite value of β , the probability to leave that site is non-zero, and, therefore, the individual will eventually pick a different choice. The time for this to happen is, however, of the order of $\exp(\beta |\phi|)$. If this time is much longer than T , we will in fact measure $h \sim 1$, even though running the trajectory for a longer time would result in $h = \mathcal{O}(1/N)$. Hence, in this case, β_c is a crossover that depends on T as $\ln(T)/(|\phi|)$.

A more interesting situation (at least from a theoretical point of view) is when $|\phi| = \infty$. As we will show below, there exists cases where β_c corresponds to a true phase transition and is independent of T (when T is large).

IV. MEAN FIELD APPROXIMATION

In order to draw further analytical features, we start by looking at a mean-field approximation. This means that we take the graph \mathcal{G} to be fully connected with $\mathcal{T}_{x_i \rightarrow x_j} = 1/(N - 1)$ and in the limit $N \rightarrow \infty$.

We now formalize the argument previously sketched. If the individual started first at a given choice corresponding to node i , then the probability $P_{\geq}(\tau)$ that he/she remains there up to a time τ is given by the product over $t \in \llbracket 0, \tau - 1 \rrbracket$ of the probabilities not to leave the site between times t and $t + 1$, $p_{\text{stay}}(t)$. Now, $p_{\text{stay}}(t) = 1 - p_{\text{leave}}(t)$, with

$$p_{\text{leave}}(t) = \sum_{j \in \partial_i} \mathcal{T}_{x_i \rightarrow x_j} \frac{1}{1 + e^{\beta[U(x_i,t) - U(x_j,t)]}}. \tag{4.1}$$

For the fully connected graph, this expression simplifies to Ref. 48 $p_{\text{stay}}(t) = [1 + e^{-\beta\Phi(t)}]^{-1}$, with $\Phi(t) = \sum_0^t \phi(s)$. It follows that:

$$P_{>}(\tau) = \prod_{t=0}^{\tau} [1 + e^{-\beta\Phi(t)}]^{-1} \approx e^{-I(\tau)}, \tag{4.2}$$

with $I(\tau) := \int_0^{\tau} dt \ln [1 + e^{-\beta\Phi(t)}]$, where we have replaced discrete sums by integrals. Equation (4.2) determines the distribution of the “trapping” time τ that the agent spends stuck on a certain choice. Its nature will entirely depend on the behavior of the integral $I(\tau)$ when $\tau \rightarrow \infty$.

A. Short term memory

Consider first the case where $\lim_{t \rightarrow \infty} \Phi(t) = |\phi| < +\infty$. Then $I(\tau) \approx \lambda\tau$ for $\tau \rightarrow \infty$, with

$$\lambda := \ln [1 + e^{-\beta|\phi|}]. \tag{4.3}$$

This means that the trapping time distribution decays exponentially fast for large τ , with an average trapping time $\langle \tau \rangle$ approximately given by $1/\lambda$. For sufficiently small λ , we recover the qualitative criterion of Sec. III by setting $T\lambda \sim 1$. But the dynamics remains ergodic when $T \rightarrow \infty$.

B. Long term memory

Suppose now that $\phi(t)$ decays sufficiently slowly for large t for $|\phi|$ to diverge. For definiteness, we will focus on power-law kernels,

$$\phi(t) = \frac{C}{(1+t)^\gamma}. \tag{4.4}$$

When $\gamma > 1$, $|\phi|$ is finite and we are back to the previous case. Hence, we restrict to $\gamma \leq 1$.

When $\gamma < 1$, one finds that $\Phi(t) \propto t^{1-\gamma}$ for large t . Hence, $I(\tau)$ converges to a finite limit I_∞ for large τ . This means that there is a finite probability $P_\infty = e^{-I_\infty}$ that the choice is made forever.

When $\gamma = 1$, $\Phi(t) \approx C \ln t$ for large t . This leads to three further sub-cases:

1. When $\beta C > 1$, $I(\tau)$ again converges to a finite limit when $\tau \rightarrow \infty$, i.e., decisions self-trap forever.
2. At the transition point, defined as $\beta_c^* = C^{-1}$, one finds that $P_{>}(\tau)$ decays as τ^{-1} , i.e., the trapping time distribution is a Zipf law, τ^{-2} . This is the marginal case that appears in several models of aging in the literature.^{34,35} For a finite observation time T , the average trapping time grows like $\ln T$.
3. When $\beta C < 1$, $I(\tau)$ behaves for large τ as τ^b/b , where $b = 1 - \beta C > 0$. The average trapping time $\langle \tau \rangle$ is thus finite. A careful analysis⁴⁹ shows that $\langle \tau \rangle$ diverges as b^{-1} when $b \rightarrow 0$, but higher moments $\langle \tau^k \rangle$ with $k > 1$ diverge much faster, as $\exp((k-1)/b)$, i.e., according to the so-called Vogel–Tamman–Fulcher law (see, e.g., Ref. 36).

Let us summarize the above results. When the kernel ϕ decays fast enough, there is a crossover regime in β between free exploration of the space of choice and trapping. The crossover value of β depends

on the observation time T and is given, using Eq. (4.3) for T large, by

$$\beta_c = \frac{\ln T}{|\phi|}. \tag{4.5}$$

When memory is long ranged, and described by a power-law kernel with decay exponent γ , there exists a genuine transition when $\gamma = 1$ between a free exploration regime and a (non-ergodic) trapped regime at a T independent value of β that we shall henceforth call β_c^* . When $\beta > \beta_c^*$ or $\gamma < 1$, there is a non-zero probability to get trapped in the same decision forever. The characteristic time for changing decision is of the same order of magnitude as T itself, a phenomenon called “aging” [see, e.g., Refs. 35 and 37 on which we will comment further below (see Fig. 4)]. Note finally that for $\gamma = 1$, our mean-field analysis predicts that while the average trapping time diverges as $(\beta_c^* - \beta)^{-1}$, all higher moments diverge much faster, as $\sim \exp(A/(\beta_c^* - \beta))$.

V. NUMERICAL RESULTS

We have conducted simulations using a long-range memory kernel given by Eq. (4.4) with $\gamma \in [1; \infty[$. For numerical convenience, we represent $\phi(t)$ as a superposition of exponentials as done in Ref. 38. We have considered a variety of different graph topologies \mathcal{G} : fully connected graphs, one-dimensional chains, and finally Watts–Strogatz small-world networks. Without loss of generality, we set again $U_0 = 1$ as this simply corresponds to a rescaling of β .

Figure 2 (left) shows the value of β_c , determined as the maximum of the variance of h , as a function of T for two different topologies (one-dimensional and fully connected) and two different values of $\gamma \in \{1, 1.5\}$. Our results show excellent qualitative agreement with the theoretical prediction for the two topologies, in particular, Eq. (4.5) in the $\gamma > 1$ case, although there is an overall factor needed to account for the one-dimensional data.

One can actually interpolate between the two situations by considering Watts–Strogatz small-world networks,³⁹ with a rewiring parameter p such that $p = 0$ corresponds to one-dimension chains and $p = 1$ to the fully connected graph. Figure 2 (right) shows the value of β_c^* as a function of the rewiring parameter p of interpolating




FIG. 2. Left: critical β_c as a function of $\ln T$ for $\gamma = 1$ and 1.5, $N = 10^5$, and different topologies. (Fc stands for fully connected, while 1D is the one-dimensional chain). Black lines correspond to the prediction of mean-field theory. Right: dependence of the critical intensity of choice β_c^* on the parameter p of Watts–Strogatz networks, for $T = 5 \times 10^3$ and $N = 2 \times 10^3$.




FIG. 3. Order parameter h as a function β for $N = 10^5$, $\gamma = 1$, and different values of T . Left: fully connected graph, for which a step function is approached as $T^{-1/3}$. Right: one-dimensional chain geometry, for which we cannot exclude that h remains a continuous function of β when $T \rightarrow \infty$.

between one-dimensional chains for $p = 0$ and the fully connected graph for $p = 1$. The parameter p , therefore, allows to interpolate between a situation where one may only do local jumps to a situation where one can go anywhere. As expected, β_c^* increases with p , as it is easier to get trapped in less connected graphs, where the same choice is revisited more often.

We now study more carefully the behavior of the order parameter h close to the transition point β_c^* , when $\gamma = 1$, both for one-dimensional chains and for the fully connected graph. We choose $N = 10^5$ henceforth, such that finite size effects are negligible in the range of T that we explore. Figure 3 suggests that as $T \rightarrow \infty$, $\langle h \rangle(\beta)$ appears to slowly converge to a step function that is zero for $\beta < \beta_c^*$ and unity when $\beta > \beta_c^*$, at least in the fully connected case where the speed of convergence is found to be $\sim T^{-1/3}$. In the one-dimensional case, one cannot exclude with the available data that this limiting function remains continuous when $T \rightarrow \infty$.

VI. AGING

Finally, let us be a little more specific about the meaning of self-trapping for finite T when $\beta > \beta_c^*$. The correct statement is that the system *ages*, in the following sense:^{37,40} assume that the agent's choice at time T is a certain x_i and ask: what is the probability $\mathcal{P}(t, T)$ that the agent has never changed his/her mind between T and a later time $T + t$? In the free exploration phase $\beta < \beta_c^*$, this probability is, for large T , independent of T : the process is time-translation invariant. In the trapped phase $\beta > \beta_c^*$, $\mathcal{P}(t, T)$ can be estimated by appropriately generalizing Eq. (4.2). The result takes the following aging form (see Fig. 4) [note that $\mathcal{P}(t, T)$ can be written in the general form advocated by Cugliandolo and Kurchan:⁴¹ $\mathcal{P}(t, T) = F(h(T)/h(T+t))$, where $F(u) = 1/u$ and $h(x)$ is an effective clock],

$$\mathcal{P}(t, T) \approx \exp\left(\frac{1}{a(T+t)^a} - \frac{1}{aT^a}\right), \quad a = \frac{\beta}{\beta_c^*} - 1 > 0, \quad (6.1)$$

where the exponent a is equal to minus the exponent b characterizing the stretched exponential distribution in the phase $\beta < \beta_c^*$ (see Sec. IV). Note that in the regime $t \ll T$, $\mathcal{P}(t, T)$ is a function of t/T^{1+a} [precisely $\mathcal{P}(t, T) \approx 1 - t/T^{1+a}$ (see dashed line in Fig. 4)],




FIG. 4. Aging function $1 - \mathcal{P}(t, T)$ as a function of $t/T^{\beta/\beta_c^*}$ in a fully connected graph with $N = 10000$, with $\beta_c^* \approx 1.5$ and $\beta = 3.5$. Equation (6.1) predicts that $\mathcal{P}(t, T) \approx 1 - t/T^{\beta/\beta_c^*}$ when $t \ll T$ (dashed line).

a regime called *super-aging*⁴² since the effective time for changing one's mind grows faster than the age T itself. This is quite interesting since we are not aware of simple models leading to such a super-linear aging behavior. Hence, memory effects of the type discussed here might very well be an interesting lead to interpret experiments that show such a super-aging behavior, such as those reported in Ref. 43.

Right at the transition point $\beta = \beta_c^*$, one finds simple aging, i.e., a scaling function of t/T ,

$$\mathcal{P}(t, T) \approx \frac{1}{1 + \frac{t}{T}}, \quad \beta = \beta_c^*. \quad (6.2)$$

When the kernel has a finite norm and leads to a crossover rather than a true transition, aging will take place whenever $T \ll e^{|\phi|}$ but revert to a normal time-translation dynamics when $T \gg e^{|\phi|}$ (see Ref. 40 for a similar situation). When $\gamma < 1$, on the contrary, relaxation is quasi-frozen for large T , in the sense that $\mathcal{P}(t, T) \approx 1 - t \exp(-\beta C T^{1-\gamma})$ when $t \ll T^\gamma$.

VII. CONCLUSION

Although quite simple, our model shows that non-trivial choice distortion effects can emerge through memory or self-reinforcing mechanisms. Our main result is that the addition of memory effects can hinder the full exploration of all choices by the agent, and it may even cause him/her to leave a substantial number of possible options totally unexplored. The emergence of aging properties also shows that including memory effects in agents' preferences can lead to non-ergodic dynamics, when ergodicity is a crucial assumption to many models in economics. Table I summarizes our results.

Several extensions can be thought of and would be a sensible way to incorporate more realism into the model. As we mentioned, we have explored here the case where the objective utility landscape $U_0(x)$ is totally flat, in a way to highlight the effects induced by memory alone. Reintroducing some heterogeneities in $U_0(x)$ would allow one to study the competition between "landscape trapping" and

TABLE I. Summary of the different dynamical regimes.

Asymptotics of $\sum_{t=0}^T \phi(t)$	$ \phi < \infty$	$\ln(T)$	$T^{1-\gamma}$, $\gamma < 1$
Asymptotics of $P_>(\tau)$	$e^{-\lambda\tau}$ with $\lambda = \ln[1 + e^{-\beta \phi }]$	<ul style="list-style-type: none"> $\beta > \beta_c^*$ (trapped regime): $P_\infty > 0$ $\beta = \beta_c^*$ (critical regime): τ^{-1} $\beta < \beta_c^*$ (ergodic regime): $\exp(-t^b/b)$, with $b = 1 - \beta/\beta_c^*$ 	$P_\infty > 0$ (trapped)
Aging $\mathcal{P}(t, T)$	Trapping and aging for $T \ll e^{\beta \phi }$	<ul style="list-style-type: none"> $\beta > \beta_c^*$ super-aging $t \sim T^{\beta/\beta_c^*}$ $\beta = \beta_c^*$ normal-aging $t \sim T$ $\beta < \beta_c^*$ equivalent to $\phi < \infty$ 	Quasi-frozen relaxation

“memory trapping,” with possibly interesting transitions between the two. One could also imagine, along the lines of Ref. 44, a situation where an agent is only sensitive to past extreme values of his/her utility.

Another direction is to introduce many agents with interactions between them, meaning that the subjective utility may also depend on what others are doing. Here again, one expects some interesting competition between herding induced condensation of choices and memory effects. Steps in this direction, with agents interacting through a market and learning through past experience, have already been taken in Refs. 45 and 46. In particular, the combined effect of imitation of the past and imitation of peers may generate *collective self-fulfilling prophecies*.

One could also explore the case where the graph \mathcal{G} defining the topology of the space of choices is itself time dependent—see Ref. 47 for a step in this direction. For example, the neighborhood of each choice could be itself affected by past choices, or some new choices, not present initially, could present themselves later in time (for example, the opening of a new restaurant).

In all these cases, the basic question is whether memory effects, habit formation, or herding completely distort the choices dictated by their objective utilities or not. Such distortions may have very significant economic consequences at the macrolevel.

From a purely theoretical point of view, revisiting reinforcement models considered in the literature²⁰ with a power-law decaying memory kernel could lead to new interesting transitions of the type discussed above. In particular, the super-aging behavior reported in the trapped phase might have applications much beyond the present setting.

ACKNOWLEDGMENTS

We warmly thank Pierre Lecointre, who contributed to the early stages of this project, as well as Alan Kirman, Jean-Pierre Nadal, and Peter Sollich for their suggestions. This research was conducted within the *Econophysics & Complex Systems* Research Chair under the egis of the *Fondation du Risque*, a joint initiative by l'École polytechnique and Capital Fund Management.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

- H. A. Simon, “A behavioral model of rational choice,” *Q. J. Econ.* **69**, 99–118 (1955).
- H. Simon, “Decision and organization. A volume in honor of Jacob Marschak,” in *Decision and Organization. A Volume in Honor of Jacob Marschak*, edited by C. B. McGuire and Roy Radner (North-Holland Pub. Co, Amsterdam, 1972), Chap. 8, pp. 161–176.
- R. Selten, “Bounded rationality,” *J. Inst. Theor. Econ.* **146**, 649–658 (1990).
- W. B. Arthur, “Inductive reasoning and bounded rationality,” *Am. Econ. Rev.* **84**, 406–411 (1994).
- G. Gigerenzer and R. Selten, *Bounded Rationality: The Adaptive Toolbox* (MIT Press, 2002).
- D. Kahneman, “A psychological perspective on economics,” *Am. Econ. Rev.* **93**, 162–168 (2003).
- D. Kahneman and R. H. Thaler, “Anomalies: Utility maximization and experienced utility,” *J. Econ. Perspect.* **20**, 221–234 (2006).
- W. A. Brock and S. N. Durlauf, “Discrete choice with social interactions,” *Rev. Econ. Stud.* **68**, 235–260 (2001).
- S. N. Durlauf and H. P. Young, *Social Dynamics* (MIT Press, 2004), Vol. 4.
- C. Borghesi and J.-P. Bouchaud, “Of songs and men: A model for multiple choice with herding,” *Qual. Quant.* **41**, 557–568 (2007).
- J.-P. Bouchaud, “Crises and collective socio-economic phenomena: Simple models and challenges,” *J. Stat. Phys.* **151**, 567–606 (2013).
- F. Guglielmo Morelli, M. Benzaquen, M. Tarzia, and J.-P. Bouchaud, “Confidence collapse in a multi-household, self-reflexive DSGE model,” *PNAS* **117**(17), 9244–9249 (2020).
- D. Kahneman, “Experienced utility and objective happiness: A moment-based approach,” in *Choices, Values, and Frames* (Cambridge University Press, 2000), pp. 673–692.
- J. Y. Campbell and J. H. Cochrane, “By force of habit: A consumption-based explanation of aggregate stock market behavior,” *J. Political Econ.* **107**, 205–251 (1999).
- A. Abel, “Asset prices under habit formation and catching up with the joneses,” *Technical Report* (1990).
- G. M. Constantinides, “Habit formation: A resolution of the equity premium puzzle,” *J. Political Econ.* **98**, 519–543 (1990).
- C. D. Carroll, J. Overland, and D. N. Weil, “Saving and growth with habit formation,” *Am. Econ. Rev.* **90**, 341–355 (2000).
- J. C. Fuhrer, “Habit formation in consumption and its implications for monetary-policy models,” *Am. Econ. Rev.* **90**, 367–390 (2000).
- R. A. Pollak, “Habit formation and dynamic demand functions,” *J. Political Econ.* **78**, 745–763 (1970).
- R. Pemantle, “A survey of random processes with reinforcement,” *Probab. Surv.* **4**, 1–79 (2007).
- D. Boyer, A. Falcón-Cortés, L. Giuggioli, and S. N. Majumdar, “Anderson-like localization transition of random walks with resetting,” *J. Stat. Mech. Theory Exp.* **2019**, 053204 (2019).
- R. L. Jack and R. J. Harris, “Giant leaps and long excursions: Fluctuation mechanisms in systems with long-range memory,” *arXiv:2003.03587* (2020).

- ²³W. Brian Arthur, Y. M. Ermoliev, and Y. M. Kaniovski, "Path-dependent processes and the emergence of macro-structure," *Eur. J. Oper. Res.* **30**, 294–303 (1987).
- ²⁴W. Brian Arthur, "Positive feedbacks in the economy," *Sci. Am.* **262**, 92–99 (1990).
- ²⁵A. P. Kirman and N. J. Vriend, "Evolving market structure: An ACE model of price dispersion and loyalty," *J. Econ. Dyn. Control* **25**, 459–502 (2001).
- ²⁶S. P. Anderson, A. de Palma, and J.-F. Thisse, *Discrete Choice Theory of Product Differentiation* (MIT Press, 1992).
- ²⁷N. Metropolis and S. Ulam, "The Monte Carlo method," *J. Am. Stat. Assoc.* **44**, 335–341 (1949).
- ²⁸W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika* **57**, 97–109 (1970).
- ²⁹B. Derrida, "Random-energy model: An exactly solvable model of disordered systems," *Phys. Rev. B* **24**, 2613–2626 (1981).
- ³⁰J.-P. Bouchaud and M. Mézard, "Universality classes for extreme-value statistics," *J. Phys. A Math. Gen.* **30**, 7997–8015 (1997).
- ³¹F. Wegner, "Inverse participation ratio in $2 + \varepsilon$ dimensions," *Z. Phys. B Condens. Matter Quanta* **36**, 209–214 (1980).
- ³²R. J. Bell and P. Dean, "Atomic vibrations in vitreous silica," *Discuss. Faraday Soc.* **50**, 55 (1970).
- ³³M. Hall and N. Tideman, "Measures of concentration," *J. Am. Stat. Assoc.* **62**, 162–168 (1967).
- ³⁴E. Bertin and J.-P. Bouchaud, "Dynamical ultrametricity in the critical trap model," *J. Phys. A Math. Gen.* **35**, 3039–3051 (2002).
- ³⁵A. Amir, Y. Oreg, and Y. Imry, "On relaxations and aging of various glasses," *Proc. Natl. Acad. Sci. U.S.A.* **109**, 1850–1855 (2012).
- ³⁶P. G. Debenedetti and F. H. Stillinger, "Supercooled liquids and the glass transition," *Nature* **410**, 259–267 (2001).
- ³⁷J. P. Bouchaud, "Weak ergodicity breaking and aging in disordered systems," *J. Phys. I* **2**, 1705–1713 (1992).
- ³⁸T. Boichud and D. Challet, "Optimal approximations of power-laws with exponentials," [arXiv:physics/0605149](https://arxiv.org/abs/physics/0605149) [physics.data-an] (2006).
- ³⁹D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature* **393**, 440–442 (1998).
- ⁴⁰C. Monthus and J.-P. Bouchaud, "Models of traps and glass phenomenology," *J. Phys. A Math. Gen.* **29**, 3847–3869 (1996).
- ⁴¹L. F. Cugliandolo and J. Kurchan, "On the out-of-equilibrium relaxation of the Sherrington–Kirkpatrick model," *J. Phys. A Math. Gen.* **27**, 5749–5772 (1994).
- ⁴²J.-P. Bouchaud, "Aging in glassy systems: New experiments, simple models, and open questions," in *Soft and Fragile Matter: Nonequilibrium Dynamics, Metastability and Flow*, edited by M. E. Cates and M. R. Evans (CRC Press, 2000), pp. 285–304.
- ⁴³L. Cipelletti, S. Manley, R. C. Ball, and D. A. Weitz, "Universal aging features in the restructuring of fractal colloidal gels," *Phys. Rev. Lett.* **84**, 2275–2278 (2000).
- ⁴⁴R. J. Harris, "Random walkers with extreme value memory: Modelling the peak-end rule," *New J. Phys.* **17**, 053049 (2015).
- ⁴⁵A. Alorić, P. Sollich, and P. McBurney, "Spontaneous segregation of agents across double auction markets," in *Lecture Notes in Economics and Mathematical Systems* (Springer International Publishing, 2014), pp. 79–90.
- ⁴⁶A. Alorić and P. Sollich, "Market fragmentation and market consolidation: Multiple steady states in systems of adaptive traders choosing where to trade," *Phys. Rev. E* **99**, 062309 (2019).
- ⁴⁷C. Cotar and V. Limic, "Attraction time for strongly reinforced walks," *Ann. Appl. Probab.* **19**, 1972–2007 (2009).
- ⁴⁸In the general case in which the agent started in a "site" different from i and then got stuck in i , one wants to replace the right-hand side by an average of the logit rule over the utility gaps $U(x_i, t) - U(x_j, t)$. However, as $N \rightarrow \infty$, it is very unlikely that a site that was previously picked is chosen again. We can thus safely replace the average gap by the gap with the base level $U_0 = 1$.
- ⁴⁹Indeed, a direct integration of the kernel yields $P_{>}(\tau) = \exp[(1 - (1 + \tau)^b)/b]$. As $b \rightarrow 0$, this function can be approximated by $P_{>}(\tau) = (1 + \tau)^{-1}$ up to a cutoff at $\tau \sim e^{1/b}$. Direct integration then gives $\int_0^{e^{1/b}} d\tau \tau^{k-1}/(1 + \tau) \approx \exp((k - 1)/b)$ for $k > 1$ and $\sim 1/b$ for $k = 1$.