Chaos

Evolutionary game dynamics of controlled and automatic decision-making

Cite as: Chaos **25**, 073120 (2015); https://doi.org/10.1063/1.4927488
Submitted: 05 July 2015 . Accepted: 16 July 2015 . Published Online: 27 July 2015

Danielle F. P. Toupo, Steven H. Strogatz, Jonathan D. Cohen, and David G. Rand

View Online Export Citation CrossMark

ARTICLES YOU MAY BE INTERESTED IN

Correlated disorder in the Kuramoto model: Effects on phase coherence, finite-size scaling, and dynamic fluctuations
Chaos: An Interdisciplinary Journal of Nonlinear Science **26**, 103105 (2016); https://doi.org/10.1063/1.4964520

Dynamics of globally coupled oscillators: Progress and perspectives
Chaos: An Interdisciplinary Journal of Nonlinear Science **25**, 097616 (2015); https://doi.org/10.1063/1.4922971

Graph partitions and cluster synchronization in networks of oscillators
Chaos: An Interdisciplinary Journal of Nonlinear Science **26**, 094821 (2016); https://doi.org/10.1063/1.4961065

© 2015 AIP Publishing LLC.

CHAOS **25**, 073120 (2015)

Evolutionary game dynamics of controlled and automatic decision-making

Danielle F. P. Toupo,[1,a)] Steven H. Strogatz,[1,b)] Jonathan D. Cohen,[2,c)] and David G. Rand[3,d)]
[1]*Center for Applied Mathematics, Cornell University, Ithaca, New York 14853, USA*
[2]*Princeton Neuroscience Institute and Department of Psychology, Princeton University, Princeton, New Jersey 08540, USA*
[3]*Department of Psychology and Department of Economics, Yale University, New Haven, Connecticut 06511, USA*

(Received 5 July 2015; accepted 16 July 2015; published online 27 July 2015)

We integrate dual-process theories of human cognition with evolutionary game theory to study the evolution of automatic and controlled decision-making processes. We introduce a model in which agents who make decisions using either automatic or controlled processing compete with each other for survival. Agents using automatic processing act quickly and so are more likely to acquire resources, but agents using controlled processing are better planners and so make more effective use of the resources they have. Using the replicator equation, we characterize the conditions under which automatic or controlled agents dominate, when coexistence is possible and when bistability occurs. We then extend the replicator equation to consider feedback between the state of the population and the environment. Under conditions in which having a greater proportion of controlled agents either enriches the environment or enhances the competitive advantage of automatic agents, we find that limit cycles can occur, leading to persistent oscillations in the population dynamics. Critically, however, these limit cycles only emerge when feedback occurs on a sufficiently long time scale. Our results shed light on the connection between evolution and human cognition and suggest necessary conditions for the rise and fall of rationality. © *2015 AIP Publishing LLC*.
[http://dx.doi.org/10.1063/1.4927488]

Dual-process theories of human cognition play a central role in the behavioral sciences. According to these theories, decisions are often made using either automatic processes that are fast and effortless but focused on the present or controlled processes that are slow and effortful but can plan for the future. Evolutionary game theory models, however, almost never consider these distinctions. Therefore, little is known about the evolutionary dynamics of automatic versus controlled processing. Here, we address this gap by introducing an analytically tractable model for the evolution of agents that use automatic or controlled processing. The agents both compete with each other and alter their shared environment. We show that under certain circumstances, automatic and controlled processing can stably coexist within the population. We also identify conditions under which limit cycles occur. In such cases, the success of controlled agents alters the environment in a way that allows automatic agents to invade and vice versa. Our results help to explain why human evolution may not necessarily be characterized by ever-increasing levels of rationality and forward-thinkingness but instead may recurrently fall prey to periods of myopia.

I. INTRODUCTION

Dual-process theories of human decision-making conceptualize decisions as arising from the interaction of (i) automatic processes that are "hardwired" and thus computationally efficient but rigid and (ii) controlled processes that are effortful but flexible.[1–7] Such a perspective has proved useful for understanding behavior across a wide range of domains and has been used heavily in fields such as neuroscience,[8,9] cognitive and social psychology,[10–16] and behavioral economics.[17–19] Yet, despite playing a key role in human evolution, the interaction (and conflict) between automatic and controlled processing has been almost entirely overlooked by evolutionary game theorists.

Controlled processing is a defining feature of human cognition, thought to underlie virtually all higher level, characteristically human cognitive functions, such as planning, problem-solving reasoning, and symbolic language—functions that, at least under some conditions, are capable of identifying and flexibly executing rational and even optimal behavior. This might be taken to suggest that evolution should favor controlled processing and that given sufficient time, control should prevail as the dominant mode of cognition. However, there is evidence that human history is characterized by cyclical dynamics that suggest a proliferation of behaviors and social structures reflective of controlled processing, only to be followed by their demise and collapse.[20,21] What might explain these historical cycles? Here, we explore the possibility that they may reflect the dynamics of interaction between automatic and controlled processing at the population level. We do so by integrating dual-process agents into an evolutionary game-theoretic framework.

We focus our investigation of automatic versus controlled processing on a particular cognitive function: intertemporal choice.[22–24] Intertemporal choice refers to decisions between options or behaviors that yield immediate

[a)]dpt35@cornell.edu
[b)]strogatz@cornell.edu
[c)]jdc@princeton.edu
[d)]david.rand@yale.edu

 © 2015 AIP Publishing LLC

reward versus those that are less rewarding in the short run but have the potential to yield greater reward in the future.

We choose to focus on intertemporal choice for three reasons. First, the prevalence of short-sighted behavior has been identified as an important contributory factor to the demise of advanced civilizations[20] and is a topic of modern concern (e.g., failures to save for retirement, overconsumption of environmental resources, and abuse of antibiotics). Second, immediacy-biased behaviors have been linked to automatic processing, while future-oriented behaviors have been linked to the engagement of controlled processing, both at the behavioral and neural levels of analysis.[8,9,17,19,25] Thus, intertemporal choice may be a useful probe for studying the consequences of interactions between automatic and controlled processing at the population level. Third, we have performed preliminary computer simulations that support this suggestion.[26] In these simulations, agents foraged for resources (e.g., food) in an environment, and either consumed found resources immediately (when using automatic processing) or according to an optimal consumption plan calculated using a complex algorithm based on past experience. Intriguingly, these simulations sometimes gave rise to evolutionary cycles in which the proportion of controlled agents in the population waxed and waned periodically. However, the complexity of the model led to analytical intractability, making it hard to understand what conditions gave rise to these cyclical dynamics and what factors were responsible for the oscillations.

A desire to understand these issues led us to the simplified model proposed in this paper. Using the replicator equation, a nonlinear dynamical system studied in evolutionary game theory,[27,28] we introduce a minimal model of dual-process agents engaged in intertemporal choice that captures the critical features of the scenario above while remaining sufficiently simple to be mathematically tractable. In doing so, we provide a formal characterization of the conditions under which cyclical dynamics emerge and the forces that drive such cycles.

II. THE MODEL

We model a world in which agents forage for goods, compete for access to these goods, and choose how to consume goods they acquire to generate fitness, with fitness being subject to diminishing marginal returns on consumption. Agents are then subject to natural selection based on their resulting fitnesses.

For simplicity, we assume there are only two types of agents, fully controlled and fully automatic, and we explore the evolution of the fraction of controlled agents, denoted as x. Automatic agents differ from controlled agents in two ways: how likely they are to acquire goods (where the speed and efficiency of automaticity is advantageous) and how they choose to consume those resources (where the rationality and planning ability of control is advantageous).

The world is parametrized by the probability ρ of finding a good (all goods are of equal size, normalized to 1 energy unit), and the competitive advantage β that automatic agents have over controlled agents in acquiring goods (where $\beta = 0$ means that both types of agents have an equal probability of acquiring goods).

A. Competitive advantage

Because automatic processing is assumed to be faster and less taxing than controlled processing, automatic agents have a competitive advantage over controlled agents when seeking to acquire goods. For example, it could be that when agents of both types simultaneously encounter a good, the automatic agent acts more quickly and snatches the good before the controlled agent can respond. Or it could be that the ponderous deliberation engaged in by controlled agents sometimes causes them to miss opportunities that an automatic agent would be more likely to exploit.

As a result, automatic agents are more likely to acquire a good in any given time period, and so the two types of agents differ in their expected waiting time between acquiring goods (i.e., the average number of time steps between acquiring one good and the next). We define the probability of acquiring a good as p_A for automatic agents and p_C for controlled agents. Thus, the average waiting time for an automatic agent τ_A is given by

$$\tau_A = \frac{1}{p_A}, \quad \text{with} \quad p_A = \rho(1 + \beta x), \tag{1}$$

while the average waiting time for a controlled agent τ_C is given by

$$\tau_C = \frac{1}{p_C}, \quad \text{with} \quad p_C = \rho\big(1 - \beta(1 - x)\big). \tag{2}$$

For $\rho > 0$ and $\beta > 0$, it is the case that $\tau_A < \tau_C$: automatic agents acquire goods more frequently than controlled agents (again, because automatic processing is faster and more efficient). Furthermore, while both p_A and p_C are increasing in x, the population average probability of finding a resource is always constant, $xp_C + (1 - x)p_A = \rho$. This is because as x increases, a greater fraction of the population is made up of controlled agents (who have a lower probability of acquiring goods than automatic agents), and this reduction in average probability exactly balances out the increase in likelihood of any individual agent acquiring a resource. Thus in the baseline model, overall resource abundance does not vary with the make-up of the population.

B. Consumption

To implement diminishing marginal returns on resource consumption, we define the fitness gained from consuming a fraction z of a good as $z/(a + z)$, where a determines the extent of diminishing marginal returns, with lower a leading to more steeply diminishing returns. Recall that goods are normalized to have size 1 when acquired.

When automatic agents acquire a good, they consume all of it immediately; hence $z = 1$, yielding a fitness benefit of $1/(a + 1)$. They then spend, on average, the next $\tau_A - 1$ time steps consuming nothing, until they again acquire a good. Therefore, the expected fitness per time step of an automatic agent is given by

$$f_A = \frac{\frac{1}{1+a}}{\tau_A} = \frac{\rho + \beta\rho x}{a + 1}, \tag{3}$$

from Eq. (1).

In contrast, controlled agents consume acquired resources more carefully: they pace their consumption, spreading it out evenly so as to obtain the maximum possible amount of fitness gain from it (because of the diminishing marginal returns on consumption, it is wasteful to consume the entire resource immediately; evenly spaced consumption results in greater fitness yield). Thus, the prudent planning of controlled agents leads them to consume $z = 1/\tau_C$ units of good in each of the τ_C time steps and thereby to gain a fitness benefit per time step of

$$f_C = \frac{\frac{1}{\tau_C}}{a + \frac{1}{\tau_C}} = \frac{\rho\big(\beta(x-1)+1\big)}{a + \rho\big(\beta(x-1)+1\big)}, \tag{4}$$

from Eq. (2).

III. EVOLUTIONARY DYNAMICS IN A CONSTANT ENVIRONMENT

Having defined the fitness of the two types of agents, we turn to evolutionary dynamics. Specifically, we ask which strategy (or combination of the two) will be favored by natural selection for different fixed values of resource availability ρ and competitive advantage of automatic agents β. We do so using the replicator equation from evolutionary game theory[27,28] to characterize how the relative fractions of controlled and automatic agents, x and $1 - x$, respectively, vary over time. The replicator equation compares the fitness of controlled agents to the population average fitness. It increases the frequency of controlled agents over time if they have higher fitness than automatic agents and decreases it if the opposite is true.

The replicator equation for our system, using (3) and (4), is given by

$$\begin{aligned}\dot{x} &= x\big(f_C - \big(xf_C + (1-x)f_A\big)\big)\\ &= (x-1)x\left(\frac{a}{a - \beta\rho + \rho + \beta\rho x} + \frac{\rho + \beta\rho x}{a+1} - 1\right).\end{aligned} \tag{5}$$

Note that we do not need a separate equation for the fraction of automatic agents because that quantity is given by $1 - x$. The long-term dynamics of (5) are characterized in Fig. 1(a), where for the sake of illustration we fix $a = 0.15$ and vary β and ρ. We see that the (β, ρ) space is subdivided into five distinct regions. We describe the dynamics within each region below.

The endpoint solutions of $x = 0$ (all automatic agents) and $x = 1$ (all controlled agents) are always fixed points regardless of β and ρ. In regions 2 and 4, these are the only fixed points. When resources are scarce and the competitive advantage of automatics is low (region 2), $x = 1$ is the global attractor and control dominates automatic processing. Conversely, when resources are plentiful and the competitive advantage of automatics is high (region 4), $x = 0$ is the global attractor and automatic processing dominates control. This is because on the one hand, automatic agents always consume $\rho\beta$ more goods on average than controlled agents in each time step (given (1) and (2)); but on the other hand,

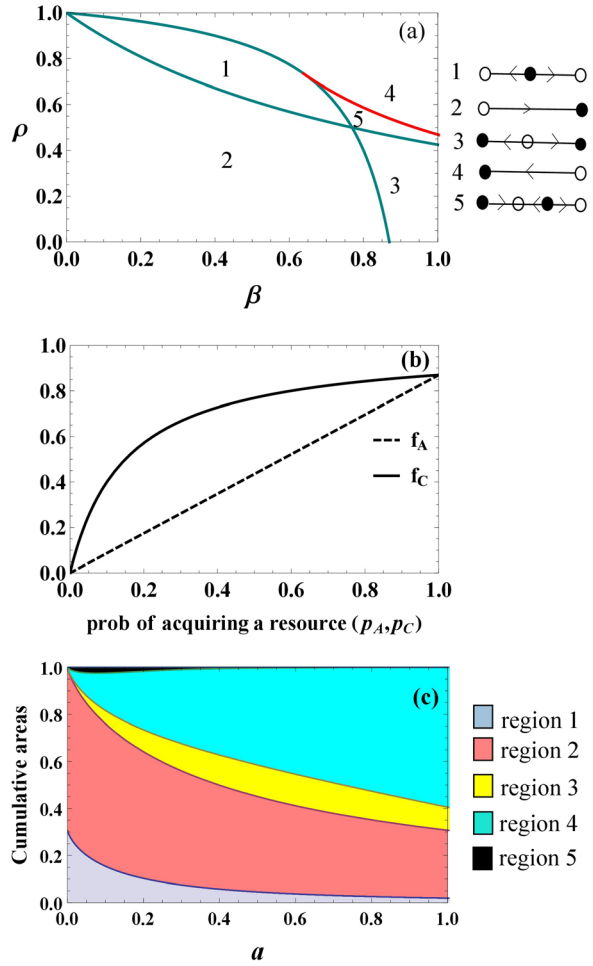

FIG. 1. Bifurcation analysis of Eq. (5). (a) Stability diagram (left) and phase portraits (right) for Eq. (5) with $a = 0.15$. Transcritical bifurcation, green curves; Saddle-node bifurcation, red curve. (b) Fitnesses f_C and f_A as functions of p_A and p_C, for $a = 0.15$. (c) Areas of regions (1)–(5) in the stability diagram, as function of a.

controlled agents make more judicious use of those resources (as controlled by a). Therefore, for a given value of a, control wins when $\rho\beta$ is sufficiently small and automaticity wins when $\rho\beta$ is large. The smaller a is (i.e., the greater the diminishing marginal returns on consumption), the larger region 2 is and the smaller region 4 is.

In the other regions, however, there can be up to two interior fixed points, in addition to these endpoint solutions. The first results from having a relatively resource-rich world with relatively little competitive advantage of automatics (regions 2 and 5). This interior fixed point is always stable and leads to coexistence of automatic and controlled processing. The second results from a relatively resource poor world in which the competitive advantage of automatics is relatively large (regions 3 and 5). This interior fixed point, by contrast, is always unstable and leads to bistability between automatic and controlled processing.

To understand why a rich world with little competitive advantage for automatics leads to coexistence while a poor world with high competitive advantage for automatics leads to bistability, we must consider how selection pressure varies based on the makeup of the population. In general, coexistence occurs when each strategy is at an advantage when it is

rare, whereas bistability occurs when each strategy is at a disadvantage when it is rare. From (1) and (2), we see that increasing the fraction of controlled agents x by a given amount also increases the probability of finding a resource for both types of agents (p_A and p_C) equally, regardless of ρ and β (although not the population-average probability of finding a resource). Therefore, what determines the dynamics when x is small versus large is how x (and the resulting increase in the probability of finding a good) translates into fitness for automatic versus controlled agents (which does depend on ρ and β). From (3) and (4), we see that f_A is linear in p_A, whereas f_C is a nonlinear function of p_C (see Fig. 1(b)). Thus, because of the concavity of f_C, an increase in the fraction of controlled agents can have different effects on the relative fitness of automatic versus controlled processing depending on ρ and β.

In a rich world (ρ large) with relatively weak competitive advantage for automatics (β small), as found in region 1, resources are common and p_C and p_A are relatively close to 1. Thus, the dynamics sit in a region where the f_C curve in Fig. 1(b) has a shallower slope than that of the linear f_A curve. Consequently, going from $x=0$ to $x=1$ leads to a bigger increase in fitness for automatic agents than controlled agents. As a result, this produces a situation in which (with the right ρ and β) control outperforms automatic near $x=0$ (when control is rare), but as x increases, the advantage of control dissipates and reverses such that automatic outperforms control near $x=1$. Thus, neither endpoint is stable, leading to coexistence.

Different dynamics occur in a poor world (ρ small) where the competitive advantage for automatics is high and β is large (region 3). Here, p_C and p_A are relatively close to 0. In this case, the slope of f_C is larger than that of f_A, and thus going from $x=0$ to $x=1$ leads to a greater increase in fitness for controlled agents than automatic agents. This produces a situation in which automatic agents outperform controlled agents near $x=0$, whereas controlled agents outperform automatic agents near $x=1$. Thus, both endpoints are stable, leading to bistability.

Finally, when both ρ and β are moderately high (region 5), the resulting long-term dynamics are a mix of regions 1 and 2, with bistability occurring between $x=0$ and a stable interior fixed point (i.e., coexistence).

Fig. 1(c) shows the areas of the five regions in Fig. 1(a) as the parameter a increases. Increasing a increases the size of the region in (β, ρ) space where automatic agents dominate (region 4) and drastically decreases the regions corresponding to bistability, coexistence, or dominance of controlled agents (regions 1, 2, 3, and 5).

IV. FEEDBACK BETWEEN THE POPULATION AND THE ENVIRONMENT

In Section III, we assumed that the environment is constant, such that the parameters ρ and β are fixed. There are many situations, however, in which the current makeup of the population can influence the environment, often with some lag.[20,21,29] Thus, in this section, we extend the model from Section III to incorporate such feedback effects. To do so, we introduce a modified version of the replicator equation that includes additional differential equations describing how the environmental parameters β and ρ vary with x, the fraction of controlled agents in the population. In Section IV A, we analyze a system in which an increase in controlled processing increases β, thus augmenting the competitive advantage of automatic agents (for example, by increasing population density). In Section IV B, we analyze a system in which an increase in controlled processing increases ρ. This scenario models a situation in which greater use of controlled processing enriches the environment and enhances resource availability for everyone, thanks (for example) to increased technological innovation leading to greater agricultural output. In Section IV C, we analyze a system with both of these features.

A. Scenario 1: Controlled processing increases competitive advantage of automaticity

Here, we consider the consequences of allowing β to positively co-vary with x. This implements a scenario in which having more controlled agents leads to greater population density and thus a larger β. The increase in population density could reflect larger population size, which results directly from the fact that populations with more controlled agents have higher average fitness. Alternatively, it could reflect an externality such as cognitive control allowing people to live more densely without violent conflict.

To link β and x, we introduce a differential equation on β that pulls its value towards the current value of x. We also incorporate the possibility of lag, specified by a parameter τ_β. This lag captures the fact that an increase in x at time t does not always have an immediate impact on β. For example, increased birth rates do not immediately lead to larger numbers of competing adults.

The new system is given by

$$\begin{aligned} \dot{x} &= x\big(f_C - \big(xf_C + (1-x)f_A\big)\big), \\ \dot{\beta} &= \frac{x-\beta}{\tau_\beta}. \end{aligned} \tag{6}$$

After insertion of (3) and (4), (6) becomes

$$\begin{aligned} \dot{x} &= (x-1)x\left(\frac{a}{a-\beta\rho+\rho+\beta\rho x} + \frac{\rho+\beta\rho x}{a+1} - 1\right), \\ \dot{\beta} &= \frac{x-\beta}{\tau_\beta}. \end{aligned} \tag{7}$$

Note that the $\dot{x}$ replicator equation is the same as it was previously, except that now β is also a variable. Also note that in equilibrium, $\beta = x$.

To illustrate how this addition of feedback between the population and the environment affects the dynamics, we begin by fixing $a=0.8$ and examining the effect of ρ and τ_β (Figs. 2(a) and 3). We find three possible types of long-term dynamics: dominance of controlled agents (region 1, no interior fixed point); coexistence of automatic and controlled agents (region 3, stable interior fixed point); and limit cycles in which both types of agents are present but their relative abundances oscillate (region 2, unstable interior fixed point).

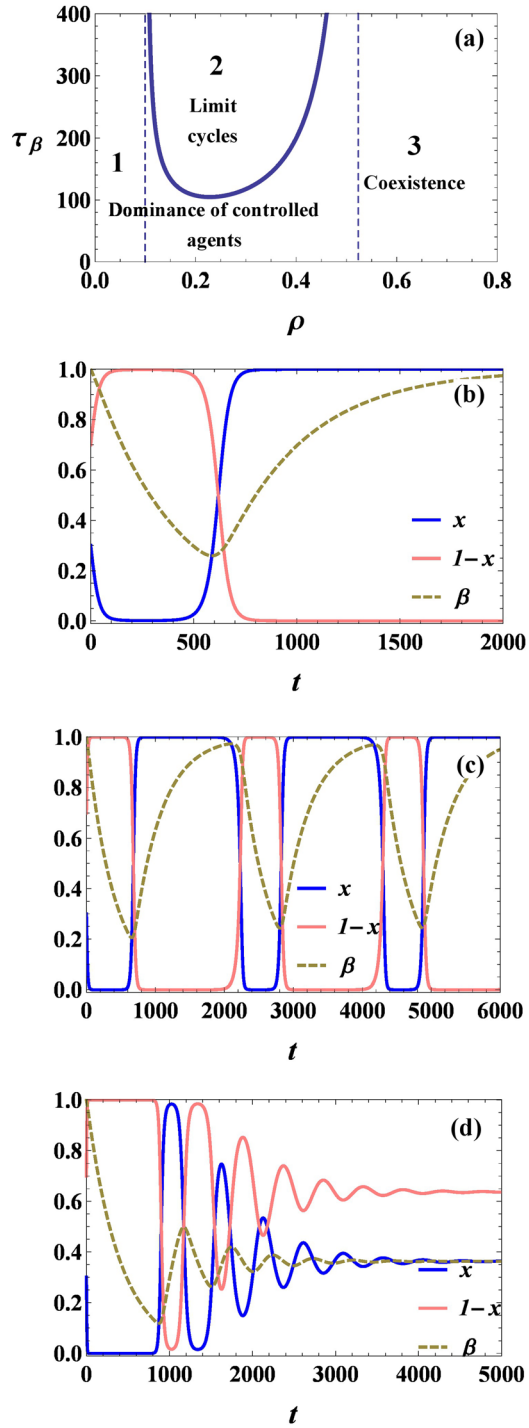

FIG. 2. Bifurcation analysis of Eq. (7) with $a=0.8$ and $\tau_\beta=400$. (a) Stability diagram. Hopf bifurcation, blue curve. (b) Time series of a typical solution in region 1 with $\rho=0.1$. (c) Time series of a typical solution in region 2 with $\rho=0.2$. (d) Time series of region 3 with $\rho=0.65$.

The parameter regime in Fig. 2(a) for which limit cycles exist is bounded by two vertical asymptotes, and within that strip, τ_β must be sufficiently large. In Fig. 2(a), $a=0.8$ and limit cycles exist if $0.1<\rho<0.52$ and $\tau_\beta>104.47$. Specifically, the limit cycles are born in a supercritical Hopf bifurcation. The equation of the Hopf bifurcation curve has been computed analytically and is too complicated to show.

We conclude this section by asking how a, the extent of diminishing marginal returns on consumption, changes the

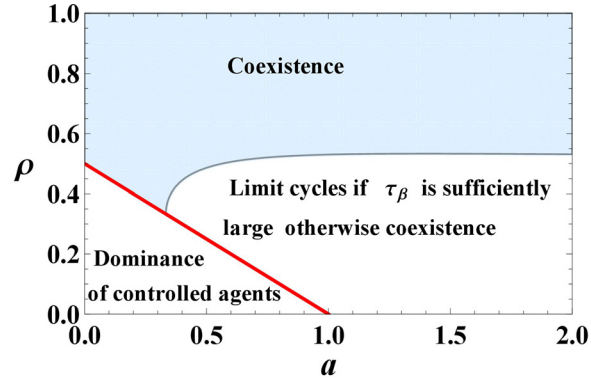

FIG. 3. Characterization of the (x, β) system from Eq. (7).

dynamics of Eq. (5) (Fig. 3). We find that only the three types of dynamics observed in Figs. 3 and 2(b)–2(d) are possible: if $\rho<(a+1)/2$, the long-term behavior is dominance of controlled agents (no interior fixed point); if $\rho>(a+1)/2$ and $\rho>\rho^*$, the long-term behavior will be coexistence; and if $\rho>(a+1)/2$ and $\rho<\rho^*$, limit cycles are possible if τ_β is sufficiently large, otherwise there will be coexistence. The curve bounding the region in (a, ρ) space where limit cycles are possible has been computed numerically.

In sum, we see that limit cycles can arise from feedback between the overall population density and the fraction of controlled agents in the population. Critically, these oscillations emerge only when the feedback is sufficiently delayed, in which case they occur over a wide range of ρ and a values.

B. Scenario 2: Controlled processing increases resource availability

Here, we leave β fixed and instead link ρ to x, using the same formulation for ρ here as for β in Scenario 1. This models a scenario in which controlled agents enrich the environment, say by creating technologies that increase resource abundance for everybody. Again, we add a lag that represents the time required for the development of such technologies and their impact on the environment to occur. This gives rise to the following system:

$$\begin{aligned}\dot{x} &= x\big(f_C - \big(xf_C + (1-x)f_A\big)\big),\\ \dot{\rho} &= \frac{x-\rho}{\tau_\rho}.\end{aligned} \tag{8}$$

After insertion of (3) and (4), Eq. (8) becomes

$$\begin{aligned}\dot{x} &= (x-1)x\left(\frac{a}{a-\beta\rho+\rho+\beta\rho x}+\frac{\rho+\beta\rho x}{a+1}-1\right),\\ \dot{\rho} &= \frac{x-\rho}{\tau_\rho}.\end{aligned} \tag{9}$$

Again, $\dot{x}$ is the same in the system without feedback, and in equilibrium $\rho=x$. For the sake of illustration, we fix $a=1.5$ and examine the dynamics as a function of β and τ_ρ (Fig. 4(a)). We find three possible types of long-term dynamics: coexistence of automatic and controlled agents (region 1, stable interior fixed point); limit cycles in which both types of

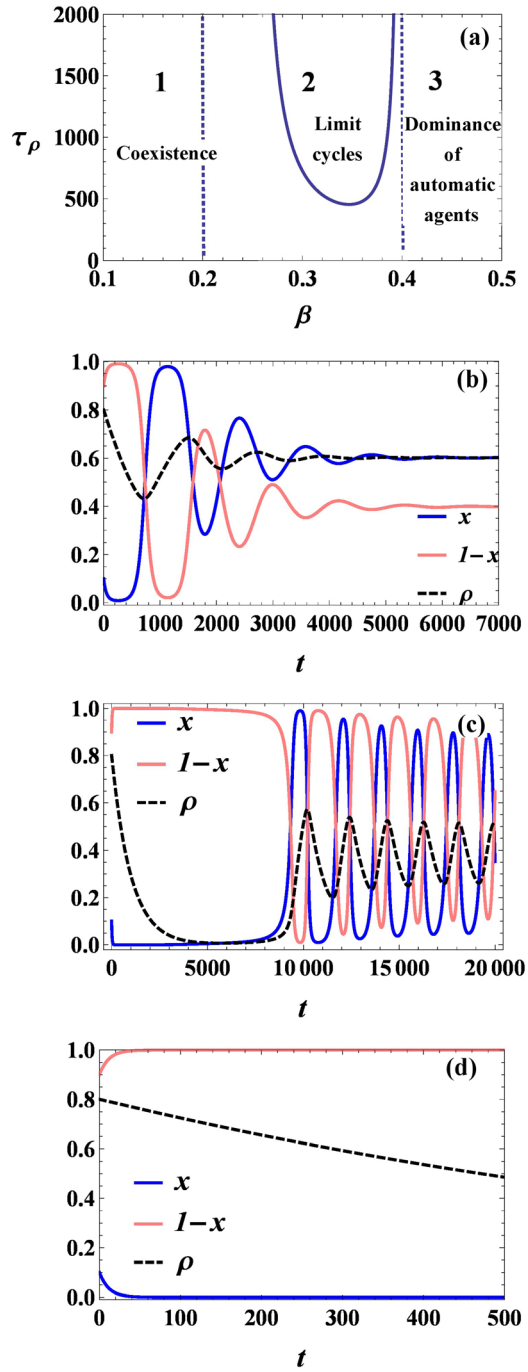

FIG. 4. Bifurcation analysis of Eq. (9) with $a = 1.5$ and $\tau_\rho = 1000$. (a) Stability diagram. Hopf bifurcation, blue curve. (b) Time series of region 1 with $\beta = 0.2$. (c) Time series of region 2 with $\beta = 0.3$. (d) Time series of region 3 with $\beta = 0.45$.

agents are present but their relative abundances oscillate (region 2, unstable interior fixed point); and dominance of automatic agents (region 3, no interior fixed point). As Fig. 4(a) indicates, limit cycles exist only if $0.249 < \rho < 0.4$ and $\tau_\rho > 446.3$. The Hopf bifurcation curve bounding the limit cycle region has been calculated analytically but is too complicated to show here. Figures 4(b)–4(d) show time series for sample trajectories from within each region.

There is an important difference between these dynamics, and the dynamics studied in Scenario 1 when x and β were positively correlated: in Scenario 1, dominance of

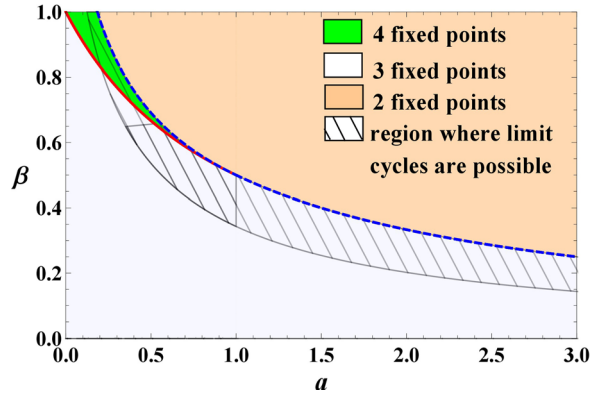

FIG. 5. Characterization of (9) with $a = 1.5$.

controlled but not automatic agents was possible, whereas here in Scenario 2, the opposite is true. Only automatic agents can dominate. Moreover, this dominance by automatic agents occurs only if $a > 1$.

Next, we ask how the dynamics of (9) depend on the parameter a, which reflects the strength of diminishing returns. We find that only the three types of long-term dynamics observed in Fig. 4(a) are possible if $a > 1$, but that more complex dynamics emerge when $a < 1$ (Figs. 5 and 6). Figure 6 characterizes the dynamics as a function of β and τ_ρ for $a = 1/2$. The long-term dynamics of (9) for $a < 1$ sometimes depend on the initial conditions, in a manner that can be summarized as follows:

- Regions 3 and 5: dominance of automatic agents.
- Region 2: either limit cycle oscillations of the two strategies or dominance of automatic agents, depending on the initial conditions.
- Region 4: either oscillations or coexistence, depending on the initial conditions.
- Regions 6 and 7: either dominance of automatic agents or coexistence, depending on the initial conditions.
- Region 8: oscillation of the two strategies, dominance of automatic agents, or coexistence, depending on the initial conditions.

In summary, adding feedback between the fraction of the population that uses controlled processing and the availability of resources can also give rise to limit cycles when the feedback is sufficiently delayed. Compared to the (x, β) system discussed in Scenario 1 (Section IV A), however, limit cycles occur over a smaller range of (β, a) combinations. Furthermore, the dynamics of the (x, ρ) system of Scenario 2 are substantially more complex than those of the (x, β) system of Scenario 1.

C. Scenario 3: Controlled processing increases both competition and resource availability

Finally, we consider the case in which x influences both β and ρ. To do so, we use a three differential equation system that includes the original replicator equation for $\dot{x}$, as well as the $\dot{\beta}$ equation from (6) and the $\dot{\rho}$ equation from (9), with the use of two different time-constants, τ_ρ and τ_β, for the two different environmental feedback equations. Thus, our system is given by

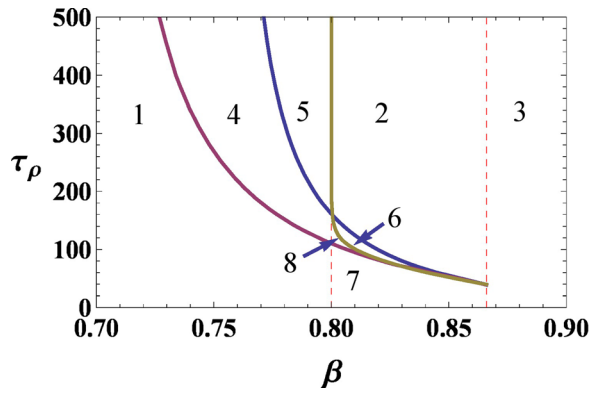

FIG. 6. Bifurcation diagram of (9) with $a = 0.5$. Hopf bifurcation, blue curve; fold bifurcation (saddle-node coalescence) of cycles, purple curve; Homoclinic bifurcation, orange curve.

$$\begin{aligned} \dot{x} &= x\big(f_C - \big(x f_C + (1-x)\big)\big), \\ \dot{\beta} &= \frac{x-\beta}{\tau_\beta}, \\ \dot{\rho} &= \frac{x-\rho}{\tau_\rho}. \end{aligned} \tag{10}$$

After insertion of (3) and (4), Eq. (10) becomes

$$\begin{aligned} \dot{x} &= (x-1)x\left(\frac{a}{a-\beta\rho+\rho+\beta\rho x} + \frac{\rho+\beta\rho x}{a+1} - 1\right), \\ \dot{\beta} &= \frac{x-\beta}{\tau_\beta}, \\ \dot{\rho} &= \frac{x-\rho}{\tau_\rho}. \end{aligned} \tag{11}$$

We perform numerical simulations to demonstrate that limit cycles can also arise in the 3D system, as shown in Fig. 7 with $a = 1.5$. We see that limit cycles are possible as long as neither τ_β nor τ_ρ are too small.

V. DISCUSSION

Here, we have introduced an analytically tractable model of the evolution of dual-process agents. Our model focuses on intertemporal choice, with agents foraging for, and competing over, goods that they consume to generate fitness. Agents that use automatic processing are at an advantage when acquiring goods because of their speed and efficiency but immediately consume any goods they acquire in short-sighted fashion. Controlled agents, conversely, engage in long-term planning and make better use of the goods they manage to acquire.

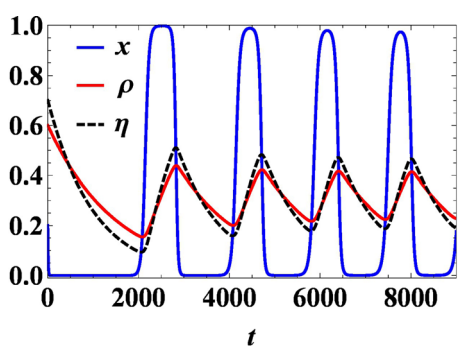

FIG. 7. Analysis of (11) with $a = 1.5$, $\tau_\rho = 1500$, $\tau_\beta = 1000$. (a) Parametric plot. (b) Time series.

Within this framework, the agents' world is parametrized by ρ, the availability of resources (defined as the average probability of finding a good per unit time), and β, the competitive advantage of automatic agents (the increased likelihood of automatic agents acquiring a good over controlled agents). Our analysis allows us to characterize which parts of the (ρ, β) parameter space lead to dominance of automatic or controlled processing, bistability, or coexistence, as well as the conditions under which limit cycles arise.

In particular, we find that natural selection favors controlled agents when ρ and β are both small (poor worlds with little competition), automatic agents when ρ and β are both large (rich worlds with substantial competition), coexistence when ρ is large and β small (rich worlds with little competition), and bistability when ρ is small and β large (poor worlds with substantial competition). Furthermore, we find that limit cycles are a robust feature of adding environmental feedback whereby a greater frequency of controlled agents leads to either higher β, higher ρ, or both. Critically, however, the feedback must be sufficiently lagged in order for limit cycles to emerge.

Thus, our analyses demonstrate the key role that feedback between the population and the environment plays in population (and ecological) dynamics. Such feedback can lead to cyclical dynamics that are otherwise impossible in a two-species competition model. Critically, environmental feedback is absent from typical evolutionary game-theoretic models, in which the game parameters are fixed, and only the population make-up varies over time.[27,28] By extending the replicator equation to include linkage between the population and one or more of the game parameters, we allow a richer range of dynamics that help to explain cyclical dynamics observed in human history.

In the interest of analytical tractability, our model makes a number of simplifying assumptions. Most importantly, we consider the limiting case of entirely automatic agents competing with entirely controlled agents. In reality, agents exist on a continuum of inclination towards automaticity versus control. We also consider a highly simplified foraging environment and a simple decision rule for controlled agents (spread consumption out evenly over the expected waiting period until the next good is acquired). We are confident, however, that these particular simplifications did not distort our results, based on our prior computer simulation work.[26] These simulations had agents that could engage in both automatic and controlled processing and examined a much more complex foraging environment. Nonetheless, our simplified model recreates the same kinds of dynamics as the more complex simulations.

The framework we introduce here can be extended in many ways to assess the impact of other simplifications and to explore other questions. For example, spatial structure could be added,[30–32] agents could differ in the extent to

which they impact the game parameters, or the game parameters could vary cyclically over time (instead of, or in addition to, variation caused by the population).[33,34] Our basic framework could also be applied to study dual-process cognition in domains beyond intertemporal choice, such as risky choice[17,35] or cooperation in social dilemmas.[14,36,37] In summary, we have introduced an evolutionary game-theoretic model of dual-process agents who make decisions using either automatic or controlled cognitive processing and who not only compete with each other but also affect their environment. Our model demonstrates how the tendency for controlled processing to enrich the environment or grow the population undermines the advantages of controlled cognition, leading to the eventual invasion of automaticity and short-sightedness. Thus, our model may shed light on historical cycles through which controlled processing, and associated phenomena such as careful planning and technological innovation, may rise and fall. The success of controlled cognition naturally leads to its own demise.

[1]D. Kahneman, *Thinking, Fast and Slow* (Farrar, Straus, and Giroux, 2011).
[2]J. D. Cohen, K. Dunbar, and J. L. McClelland, Psychol. Rev. **97**, 332–361 (1990).
[3]E. K. Miller and J. D. Cohen, Annu. Rev. Neurosci. **24**, 167–202 (2001).
[4]M. I. Posner and C. R. R. Snyder, *Attention and cognitive control in information processing and cognition: The Loyola symposium*, edited by R. L. Solso (L Erlbaum Associates, Hillsdael, NJ, 1975), pp. 55–85.
[5]W. Schneider and R. M. Shiffrin, Psychol. Rev. **84**, 1–66 (1977).
[6]R. M. Shiffrin and W. Schneider, Psychol. Rev. **84**, 127–190 (1977).
[7]H. C. Barrett and R. Kurzban, Psychol. Rev. **113**, 628–647 (2006).
[8]S. M. McClure, D. I. Laibson, G. Loewenstein, and J. D. Cohen, Science **306**, 503–507 (2004).
[9]T. A. Hare, C. F. Camerer, and A. Rangel, Science **324**, 646–648 (2009).
[10]J. S. B. Evans, Trends Cognit. Sci. **7**, 454–459 (2003).
[11]J. S. B. Evans and K. E. Stanovich, Perspect. Psychol. Sci. **8**, 223–241 (2013).
[12]K. E. Stanovich and R. F. West, Behav. Brain Sci. **23**, 645–665 (2000).
[13]A. Tversky and D. Kahneman, Psychol. Rev. **90**, 293–315 (1983).
[14]D. G. Rand, J. D. Greene, and M. A. Nowak, Nature **489**, 427 (2012).
[15]M. J. Crockett, Trends Cognit. Sci. **17**, 363 (2013).
[16]F. Cushman, Pers. Soc. Psychol. Rev. **17**, 273 (2013).
[17]D. Fudenberg and D. K. Levine, Am. Econ. Rev. **96**, 1449–1475 (2006).
[18]D. Kahneman, Am. Econ. Rev. **93**, 1449–1475 (2003).
[19]R. H. Thaler and H. M. Shefrin, J. Political Econ. **89**, 392–406 (1981).
[20]J. Diamond, *Collapse: How Societies Choose to Fail or Succeed* (Penguin, 2005).
[21]P. J. Richerson, R. Boyd, and R. L. Bettinger, Human Biol. **81**, 211–235 (2009).
[22]G. Ainslie, Psychol. Bull. **82**, 463–496 (1975).
[23]D. Laibson, Q. J. Econ. **112**, 443–478 (1997).
[24]R. Thaler, Econ. Lett. **8**, 201–207 (1981).
[25]A. Ward and T. Mann, J. Personal. Soc. Psychol. **78**, 753–763 (2000).
[26]D. Tomlin, D. G. Rand, E. A. Ludvig, and J. D. Cohen, Sci. Rep. **5**, 1–11 (2015).
[27]J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics* (Cambridge University Press, 1998).
[28]M. A. Nowak, *Evolutionary Dynamics* (Belknap Press, 2006).
[29]J. D. Cohen, J. Econ. Perspect. **19**, 3–24 (2005).
[30]R. Durrett and S. Levin, Theor. Popul. Biol. **46**, 363–394 (1994).
[31]M. A. Nowak, C. E. Tarnita, and T. Antal, Philos. Trans. R. Soc. B: Biol. Sci. **365**, 19–30 (2010).
[32]M. Perc, J. Gómez-Gardeñes, A. Szolnoki, L. M. Floría, and Y. Moreno, J. R. Soc. Interface **10**, 20120997 (2013).
[33]R. H. Rand, M. Yazhbin, and D. G. Rand, Commun. Nonlinear Sci. Numer. Simul. **16**, 3887–3895 (2011).
[34]R. E. Ruelas, D. G. Rand, and R. H. Rand, Proc. Inst. Mech. Eng., Part C **226**, 1912–1920 (2012).
[35]H. B. Zur and S. J. Breznitz, Acta Psychol. **47**, 89 (1981).
[36]D. G. Rand and M. A. Nowak, Trends Cognit. Sci. **17**, 413 (2013).
[37]D. G. Rand, A. Peysakhovich, G. T. Kraft-Todd, G. E. Newman, O. Wurzbacher, M. A. Nowak, and J. D. Greene, Nat. Commun. **5**, 1–12 (2014).