

SCIENTIFIC REPORTS



OPEN

The evolution and devolution of cognitive control: The costs of deliberation in a competitive world

Damon Tomlin¹, David G. Rand², Elliot A. Ludvig^{3,4} & Jonathan D. Cohen^{3,5}

Received: 25 March 2015

Accepted: 17 April 2015

Published: 16 June 2015

Dual-system theories of human cognition, under which fast automatic processes can complement or compete with slower deliberative processes, have not typically been incorporated into larger scale population models used in evolutionary biology, macroeconomics, or sociology. However, doing so may reveal important phenomena at the population level. Here, we introduce a novel model of the evolution of dual-system agents using a resource-consumption paradigm. By simulating agents with the capacity for both automatic and controlled processing, we illustrate how controlled processing may not always be selected over rigid, but rapid, automatic processing. Furthermore, even when controlled processing *is* advantageous, frequency-dependent effects may exist whereby the spread of control within the population undermines this advantage. As a result, the level of controlled processing in the population can oscillate persistently, or even go extinct in the long run. Our model illustrates how dual-system psychology can be incorporated into population-level evolutionary models, and how such a framework can be used to examine the dynamics of interaction between automatic and controlled processing that transpire over an evolutionary time scale.

Humans have a remarkable capacity for rational, deliberative thought. This can be especially important in considering the benefits of long-term goals and planning accordingly. Why is it, then, that we often appear to behave irrationally, succumbing to immediate temptations and acting against our own long-term interests? Part of the answer may be found in dual-system theories of cognition^{1–6}. These theories posit two types of processes: automatic ones that produce rapid, stereotyped responses with little effort and controlled ones that are slower and more effortful, but also more flexible. Dual-system theories have received considerable support in psychology^{7–10}, economics^{11–13} and neuroscience^{14,15}. Understanding the interaction of these two systems may be essential for understanding important social and ecological problems that we confront in the modern world, such as shortfalls in retirement savings, increases in the prevalence of obesity and drug addiction, depletion of the environment, and other seemingly irrational behavior^{16,17}. In this paper, we use an evolutionary model to begin to generate such an understanding.

Automatic processes can be viewed as responses optimized to specific internal or environmental signals^{18–20} that have emerged either over the course of evolution (e.g., “hardwired” reflexes) or as a consequence of extensive training (e.g., athletic skills or reading). Though these responses may be optimal within the context in which they developed, they are inflexible (i.e., context insensitive) and, if elicited under circumstances other than those in which they were developed (such as the current environment faced by most humans), may prove to be suboptimal^{21,22}.

¹Department of Psychology University of Colorado, Colorado Springs Colorado Springs, CO 80918, USA.

²Department of Psychology, Department of Economics, School of Management Yale University, New Haven, CT 06520, USA. ³Princeton Neuroscience Institute Princeton University, Princeton, NJ 08540, USA. ⁴Department of Psychology University of Warwick, Coventry, CV4 7AL, UK. ⁵Department of Psychology Princeton University, Princeton, NJ 08540, USA.

Correspondence and requests for materials should be addressed to J.D.C. (email: jdc@princeton.edu)

Intriguingly, perhaps the factor most responsible for the aforementioned change in the human environment is the emergence of the human capacity for controlled processing. Controlled processing (often referred to as “cognitive control” or “executive function”) and the flexibility it affords is fundamental to all of the faculties considered to be characteristically human — including reasoning, problem solving, language, and planning. These faculties, in turn, have given rise to an explosion of technologies that have dramatically changed our environment. However, we still carry with us those evolutionarily older automatic processes that may, in some circumstances, favor different behaviors than controlled processing (for more on the distinction between automatic and controlled processing, see Supplementary Text S1).

Why has the human brain preserved automatic responses when the ability to respond flexibly would seem to confer considerable adaptive advantage? Upon first inspection, it may seem that controlled processing should dominate automatic processing. While deliberation may be effortful and take time^{3,7,23}, it has the considerable advantage of allowing the agent to respond effectively to its current circumstances and anticipate future ones. Thus, agents with the capacity for flexible responding should perform more favorably and, over the course of evolution, come to dominate the population. There are several reasons, however, why this may not be an inexorable outcome. For example, evolution toward exclusively controlled processing may be constrained by architectural factors that preserve structures underlying automatic processing²⁴.

We propose an alternative: that interactions between controlled and automatic processing, and their interaction with the environment, may constrain the prevalence of controlled processing. More specifically, we consider the possibility that these interactions produce evolutionary dynamics in which (a) automatic processing is (or may become) more adaptive, or (b) that there is no stable equilibrium, but rather a continual oscillation in the dominance of the two types of processing. These dynamics may occur because the very success of controlled processing may lead to its downfall: the benefits that control allows may support either overgrowth of the population, or the invasion of agents that rely entirely on automatic processing, a type of “free loader” effect.

In the work presented here, we investigate the interplay between simplified forms of automatic and controlled processing in a particular setting: an environment in which agents in a population gather a resource (e.g., food) in order to survive and reproduce. We represent automatic and controlled processing in terms of two differing consumption policies: one that is rigid and has a strong immediacy bias (automatic), and the other that is flexible, and adapts to the environment in a future-oriented manner (controlled). While this domain — often referred to as intertemporal choice — is only one of many in which dual-system accounts of cognition are relevant, there are several reasons why it is an attractive starting point for studying interactions among dual-system agents at the population level.

First, it is a domain in which automatic and controlled processing make distinguishable contributions to behavior and there is growing evidence that automatic processing strongly discounts the future relative to controlled processing^{11,13,15,25,26}. For example, interfering with controlled processing biases decision-makers toward options that are immediately gratifying but less remunerative^{25,27} or even possess negative consequences in the long term²⁸. This is supported by evidence that the strong bias toward immediate rewards observed in human discounting behavior^{29–31} relies on evolutionarily older brain structures^{7,14}, while future-oriented behavior relies on controlled processes that engage more recently evolved structures such as the prefrontal cortex^{15,21,32–35}. Furthermore, theorists have provided formal arguments that the observation of preference reversals (e.g., choosing a deferred reward over an earlier but lesser one when both are in the future, but the lesser one as they draw closer in time^{36,37}) and pre-commitment (choosing to make an immediately desirable resource unavailable so as to ensure the receipt of a more valuable one later^{38,39}) can be rationally explained only in terms of competing systems of valuation⁴⁰, a formal observation that is consistent with dual-system theories of cognition. Finally, intertemporal choice is a domain in which optimal performance can be precisely defined and operationalized in a relatively straightforward manner, thus permitting the design of agents that behave in a well specified and interpretable manner.

While there is growing consensus that controlled and automatic processing may compete within the individual to determine the outcome of intertemporal choice, there has been almost no work examining how agents with different biases may interact at the population level (though this question has begun to be examined in other domains⁴¹). For simplicity, consider a population in which there are two types of agents: one relying entirely on automatic processing (“automatic agents”) and another relying entirely on controlled processing (“controlled agents”). Assume that automatic agents do not have the flexible cognition necessary to consider the future, and thus exhibit a rigid policy with a strong immediacy bias^{13,40} – for example, they may consume all (or some fixed portion) of the resources that they encounter, irrespective of the availability of resources in the environment. In contrast, allow that controlled agents can deliberate, deciding how much of a resource to consume immediately, and how much to store for the future. The key distinction is *flexibility* – even if automatic agents did not immediately consume all resources but instead consumed them in some prescribed manner (such as genetically hardcoded manner in which squirrels save for the winter in a specific way), and even if this level could evolve over generations to adapt to new environments, control would still have an advantage because of the agent’s ability to adapt its behavior to current circumstances (e.g., current level of sustenance, estimate of the likelihood of finding more resources, and the amount of resource it has already stored). That is, the

flexibility of controlled processing confers the ability to adapt behavior on an individual timescale, rather than an evolutionary one.

However, suppose that such flexibility comes at a cost, for example the extra time it takes to deliberate as compared to the implementation of a fixed and fast consumption policy^{7,23}. Indeed, speed of processing is one of the defining features of automaticity^{3,5} and may provide an advantage to automatic agents in some contexts. For example, if an automatic must compete to acquire a good (e.g., with a controlled agent), use of a predefined policy may allow the automatic agent to act faster to consume the resource — while controlled agents stop to think about what would be optimal. We point to speed as one advantage of automatic processing, but there may be others — for example, efficiency (i.e., automatic behaviors may require less energy) and reliability. This tension between controlled and automatic processing may be a general property of intelligent systems — for example, it parallels a similar tension in computational systems between interpreted procedures (such as high level programming languages, that are flexible and adaptable), and compiled procedures (such as drivers, which are fast but limited in scope). Each has its advantages. What is less clear is how agents with these differing attributes interact.

Would the speed of automaticity outweigh the flexibility of controlled processing? How would factors such as the prevalence of resources or other agents affect the relative advantage of each policy and competition at the population level? How might secondary effects of controlled processing (such as increasing population size or the availability of resources) interact with these factors, and what would be the dynamics of these interactions? The answers to these questions are not obvious. Here, we introduce a theoretical framework for examining the dynamics of interactions among dual-process agents at the evolutionary timescale, and apply it to shed light on these questions.

Methods

Population model. In the conceptual example above, there were two distinct types of agents: automatic and controlled. In our simulations, we considered a more realistic case involving actual *dual*-process agents that used both kinds of processing: each agent had some probability of overriding its automatic response using controlled processing. For simplicity, the model treated this distinction between automatic and controlled processing as dichotomous, holding aside the issues that cognitive processes may exist along a continuum⁴² and may even be relative: that is, a process may be “automatic” or “controlled” depending on the process(es) with which it competes¹. Also for simplicity, we treated the choice between automatic and controlled processing at each time point as based on a fixed probability, holding aside contextual and strategic factors that might influence this decision. We return to a consideration of these interesting and potentially important issues in the Discussion. Agents encountered resources over time, chose how to consume versus store those resources, and competed with one another for access when they encountered the same resource simultaneously.

Energy and fitness. We examined how the probability of using automatic versus controlled processing evolved in a population of dual-system agents as they encountered and consumed resources. Agents inhabited an environment in which they had a specified probability of finding a resource in each time period, and a specified probability of encountering another agent in each time step (we varied the number of agents to study the effects of competition; see Supplementary Text S2 for additional details). Each agent had an *energy level* E that varied over time and directly determined reproductive success (fitness). On each time step, E was decreased by a constant amount, representing the energy required to subsist. To offset this drain on their stored energy, agents consumed a single type of resource that was encountered on each time step with fixed probability R (*environmental “richness”*). When acquired, resources could either be consumed to increase E or, when acting in a controlled manner, stored to increase the agent’s reserve of resources, S (see below), for use in the future. Consuming resources increased E with diminishing effect, implemented as a concave utility function whereby consumption produced the greatest increases in E when it was lowest and lesser increases as E increased (see Supplementary Text S3 and Figure S1). Thus, agents implemented three important biological constraints: 1) living organisms continuously expend energy for survival (drain), 2) consuming exogenous resources yields energy (resource-dependence), and 3) consumption is associated with satiety (diminishing returns).

Dual-process agents. We implemented the distinction between automatic and controlled processing in two simple ways, intended to capture the most fundamental features of this distinction. First, automatic processing always resulted in the full consumption of an encountered resource, while controlled processing was more flexible, using a consumption policy that was sensitive to current environmental conditions and the future value of the resource. (We also performed simulations in which the automatic policy changed over evolutionary time, adapting a fixed level of consumption from generation to generation; see Supplementary Text S8). Second, when an agent acting automatically and an agent acting in a controlled fashion encountered a resource at the same time, the agent using automatic processing acquired the resource (and the controlled agent got nothing), reflecting the greater speed (or efficiency) of automatic processing. We chose this implementation so as to strike a balance between capturing the most important and widely agreed-upon features of automatic and controlled processing, while keeping the model from becoming excessively complex^{3–5,9,23}. For example, a more complex implementation of controlled processing might take account of the nature of an encounter (e.g., whether the competitor

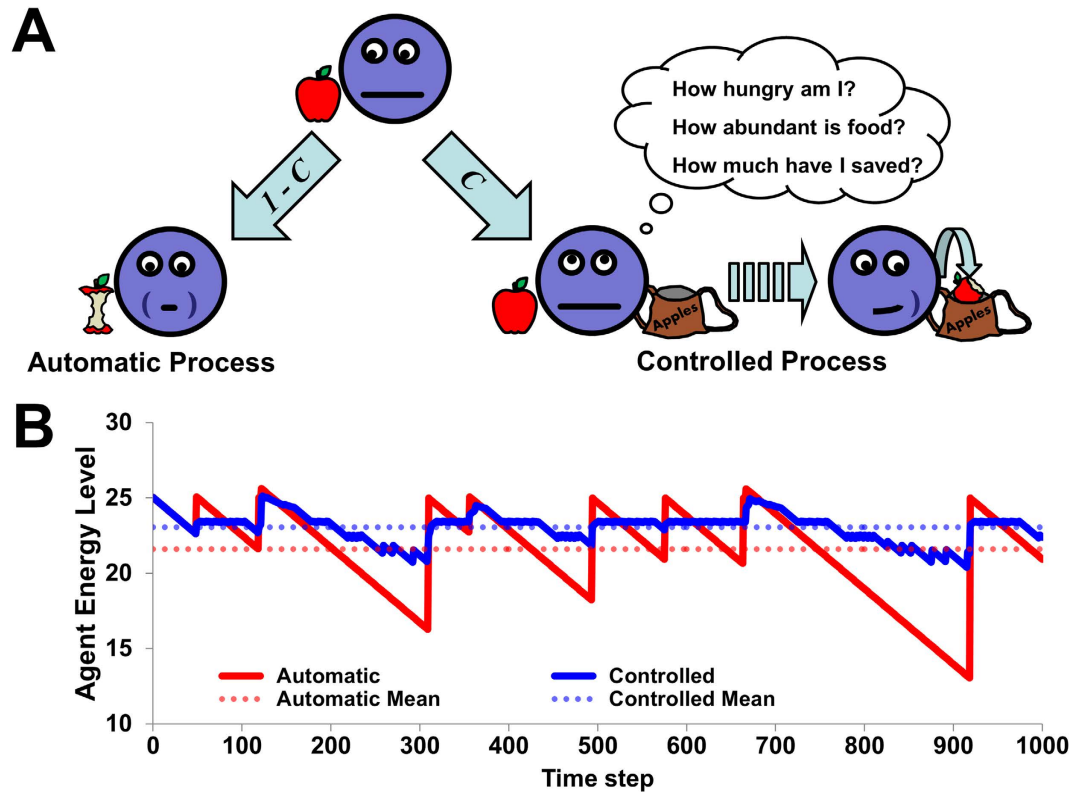


Figure 1. Consumption behavior and performance for automatic and controlled processing. (A)

An automatic agent immediately and fully consumes any resource it encounters with probability $1 - C$. A controlled agent uses its current energy level, how much it has stored, and the likelihood of finding resources, to determine the optimal amount to consume now versus store for later consumption (with probability C). However, controlled agents lose resources to the faster automatic agents when they are present simultaneously. (B) Sample energy levels over time for two agents: one using only automatic processing (red line), and the other using only controlled processing (blue line). The automatic agent experiences large, transient increases to its energy level, while the controlled agent maintains a higher mean energy level by storing resources and using them more efficiently. Energy levels are shown relative to worst possible mean performance (no consumed resources over the entire simulation).

is more likely to act in an automatic or controlled manner). While such considerations are clearly an important direction for future research (as we consider in the Discussion), even the simple model we describe exhibits complex dynamics, and understanding these may provide a valuable foundation for the examination of more complex models in the future.

We assumed that automatic processing was the default behavior, but that each agent had a stable probability C of over-riding this default and acting in a controlled manner. Thus, in a given time step, an agent acted in either a controlled or automatic manner (probabilities C and $1 - C$, respectively) regardless of whether the agent encountered a resource or had to compete for that resource with another agent. If the agent acted automatically and encountered a resource, it consumed the entire resource and its energy level E was updated accordingly (Fig. 1A). If it did not encounter a resource, its energy level was simply decremented by the energy drain. In contrast, if the agent acted in a controlled manner and encountered a resource, it estimated an optimal amount x to consume from the resources available (the resource it just encountered plus those it had already stored, if any) through a process described below. By consuming the amount x , the agent increased E according to the concave utility function. Any unconsumed resources were stored, and the value of the agent's energy reserve, S , was updated accordingly. If the agent acted in a controlled manner but did not encounter a resource, its energy level was decremented by the energy drain, the agent then chose an amount x to consume from its current store, and S was adjusted accordingly. Because agents acting automatically could not access the store, controlled processing in our model involved an idealized form of pre-commitment, in which stored resources could be accessed online in an optimal way and were unavailable when using automatic processing (for a relaxation of such constraints on automatic processing, see Supplementary Text S8, Figure S2 and Figure S3).

Competition over resources. As we discuss in greater detail below, controlled behavior was optimal with respect to consumption. As noted above, however, this came at a cost: in direct competition for a

resource, agents acting in a controlled manner lost to those acting automatically. In such instances, the automatic agent fully consumed the resource, while the agent using controlled processing got nothing but could still consume previously stored resources (in an alternate version of the model, we varied the probability of the outcome of the competition; as long as the probability that the automatic agent won was above 0.5, the results were qualitatively similar to those described here). If both agents used the same type of processing, then the resource was allocated randomly to one of them, and they applied their respective consumption policies. The probability of two agents simultaneously encountering a resource (and therefore competing over it) in a given time step increased with the population size (see Supplementary Text S2 for details).

Optimal consumption under controlled processing. Because consumption was subject to saturation, the benefit from each additional unit consumed diminished as an agent's *energy level* E increased. Optimal consumption therefore demanded that the benefit of consumption be weighed against the future value of the resource if it were stored and consumed later. If the probability of acquiring resources was low (i.e., the resource was scarce and/or competition was high), it would be wiser to save some of the resource for the future, when E was likely to be lower and the benefit of consumption higher. Optimal consumption thus depended on the current value of E , an estimate of the future availability of resources in the environment, and the amount of resources already stored.

The estimate of resource availability was experience-based, and thus sensitive to the effect of competition (because resources collected by other agents were not a part of the estimating agent's experience). When agents used controlled processing, they took these considerations into account and balanced immediate consumption with storage so as to maximize E and thus fitness (see Supplementary Text S4 for details regarding this calculation).

Figure 1B depicts the energy levels of two different agents over time: one agent that always behaved automatically (red line) and another that always used controlled processing (blue line), each encountering an identical sequence of resources. The controlled agent exhibited both less volatility and a consistently higher mean E by storing an appropriate portion of its resources for later consumption. However, this superiority depended on the size of the population and the nature of the environment.

Computation of agent fitness. Given the framework described above, an agent's fitness was determined by 1) its level of control, 2) the probability that the agent would have access to a resource when it was in an automatic state, and 3) the probability that the agent would have access to a resource when it was in a controlled state. These latter two probabilities can be readily calculated from the *environmental richness* R , the *population size* N , and the distribution of *controlled processing* C across those agents. The two probabilities, weighted by the agent's own value of C , determined the final probability that the agent would acquire a resource on any single time step. In each simulation, agents were given access to resources stochastically according to this single probability; for time steps on which the agent acquired a resource, the agent's *energy level* E was updated according to its consumption policy (which was determined according to the probability C on each time step), and the next time step began. Prior to the evolutionary simulations (see below) we precomputed agent fitness, defined as mean energy level E , for a uniform sample of the possible values of the three determinants of agent fitness listed above through 30,000 simulations of 1,000 time steps each. In the evolutionary simulations, R , N and C were used to index this precomputed table of fitnesses, and the corresponding mean E was used for the agent's fitness (see Supplementary Text S5 for details regarding these calculations).

Evolutionary simulations. The population of N agents, each of which had its own value of *controlled processing* C , evolved according to the Wright-Fisher process⁴⁵. In each generation, the population was reconstituted by sampling proportional to fitness from the previous generation, with replacement. The simulations included local mutation: with 5% probability, each new agent mutated, randomly adding or subtracting 0.02 from the value of C inherited from its parent (with C constrained to be between 0 and 1). Because we were interested in the spread of controlled processing within the population, simulations were always initialized with a homogeneous population of agents, each with $C=0$. Thus, for each population of N agents in an environment with richness R , we computed distributions of control levels as they evolved across generations.

Interaction between population size N , environmental richness R , and controlled processing C . We first examined how a fixed *population size* N and *environmental richness* R interacted with the distribution of values for *controlled processing* C in the population to influence the evolution of C . Together, N and R determined the likelihood of competition between agents over resources, which increased for higher values of both (see Results). The goal of later simulations was to characterize the dynamics in scenarios in which either N or R varied as a function of E and the propensity for controlled processing C . In *variable population size* simulations, R was held constant while increases in the mean E of the population increased N . In *variable richness* simulations, N was held constant while increases in the mean E of the population, coupled with higher frequencies of controlled processing, increased R . The latter explored the idea that surplus energy can be exploited by controlled processing (e.g., via technological innovation) to produce surplus resources.

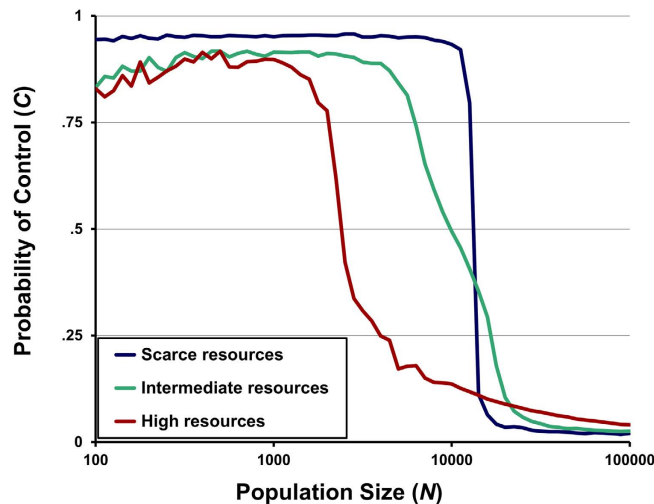


Figure 2. Equilibria depend on both *population size* N and *environmental richness* R . The x-axis represents the *population size* N , while the y-axis represents the propensity for *controlled processing* C . Each line corresponds to a different value of *environmental richness* R , with each data point indicating the mean value of C at the end of 100,000 generations for a given combination of N and R . Values of R were .005, .025, and .125 for “scarce,” “intermediate,” and “high” resources, respectively.

Variable population size simulations. Here, we considered the effect of population growth by allowing *population size* N to positively co-vary with population fitness. The population was initialized with a nominal population size N_0 . In each subsequent generation, if the mean *energy level* E of the population was above a specified threshold T_N by some margin (i.e., above $T_N + \epsilon$), then N increased by a small fixed number of individuals, whereas if the mean E was below T_N by the same margin (i.e., below $T_N - \epsilon$), N decreased by that number of individuals. If the mean E was between $T_N - \epsilon$ and $T_N + \epsilon$, N remained constant, thus allowing for a population of stable size (rather than one that continually oscillated between two sizes). Thus, the threshold T_N parameterized the difficulty of maintaining a population’s size, and simulations explored the effects of different values of T_N and initial population size N_0 (see Supplementary Text S6 for additional details).

Variable richness simulations. Here, we considered the effect of innovation resulting from the interaction between surplus fitness (i.e., energy in excess of that needed for survival) and controlled processing (which could exploit surplus fitness to enrich the environment). For these simulations, the environment was initialized with nominal richness R_0 , representing the rate at which the environment yielded resources in the absence of innovation. In each generation, if the product of the population’s mean *energy level* E and its mean value for *controlled processing* C was above a specified threshold $T_R + \epsilon$, *environmental richness* R increased by a fixed amount. This method therefore represented a simple version of the above concept – that innovations (such as environmental improvements) require the combination of deliberation and energy in order to be conceived and implemented. If $E \cdot C$ was below $T_R - \epsilon$, R decreased by the same amount, but it was never allowed to fall below R_0 (i.e., the baseline richness of the environment). If $E \cdot C$ was between $T_R - \epsilon$ and $T_R + \epsilon$, R remained constant, thereby allowing for a population with a stable level of environmental richness. Thus, the threshold T_R parameterized the difficulty of modifying the environment, and simulations explored the effects of different values for T_R and R_0 (see Supplementary Text S7 for additional details).

Results

Evolutionary outcomes for fixed combinations of population size N and environmental richness R . Figure 2 shows the results of initial simulations in which *population size* N and *environmental richness* R had fixed values. Higher values of N favored automatic processing because a larger N yielded a higher probability of competition (in which case automatic agents won the resource). Higher values of R also favored automatic processing because greater resource availability decreased the benefit of storing resources for future consumption.

These simulations also showed that N and R interacted to determine the success of controlled versus automatic processing: lower levels of R were associated with a greater range of N for which controlled processing was an equilibrium. Importantly, at higher values of R , there was also a range of values for N at which equilibria were reached with an intermediate value of *controlled processing* C — that is, in which dual-processing was a stable strategy, with agents sometimes acting automatically and sometimes acting in a controlled fashion (with a fixed probability).

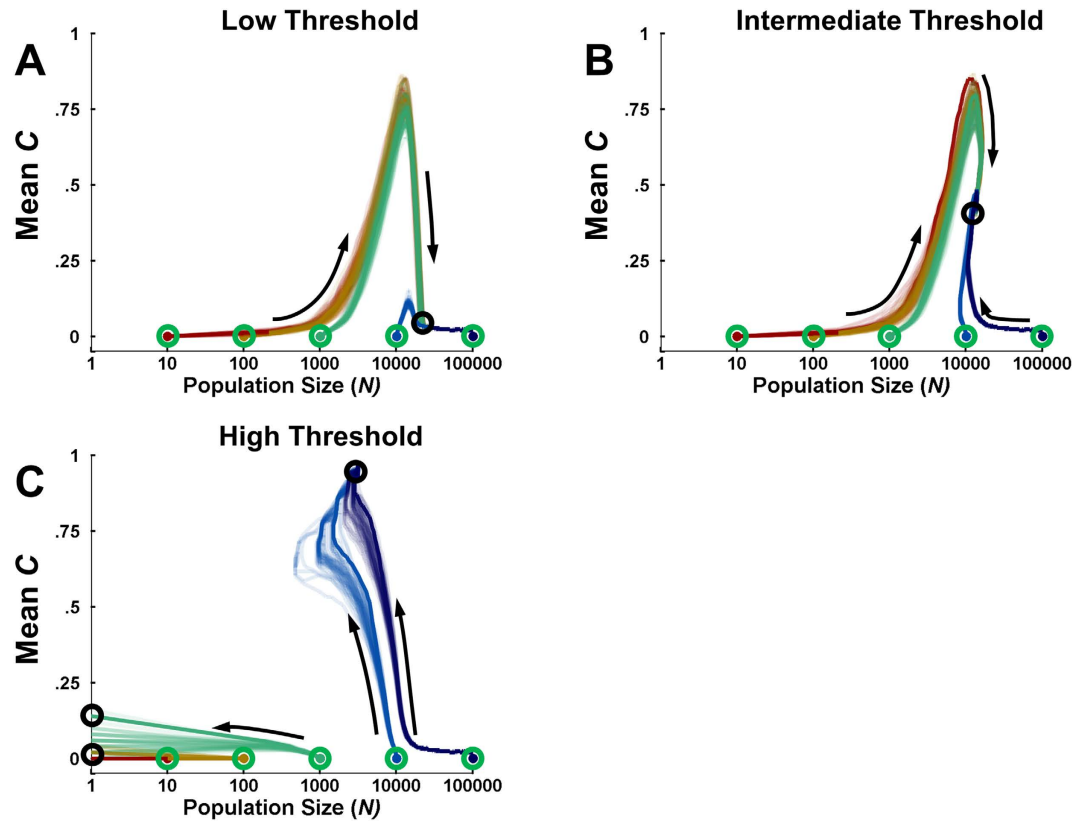


Figure 3. Variable population size simulations. Each panel shows evolutionary trajectories, starting with a value of zero for *controlled processing C* but with incrementally larger initial population sizes (x-axis). For each set of trajectories, a single exemplar is shown with a dark line, with dark circles indicating the final disposition of these exemplars. (A) When the threshold T_N was small and maintaining *population size N* was relatively easy, small populations evolved *controlled processing* until the increased population size made automatic processing more advantageous. (B) For an intermediate value of T_N , populations reached a stable size only after evolving a moderate degree of *C*. (C) For high values of T_N , populations required a very high fitness in order to maintain their current size, and only those that begin with high *N* evolved sufficient control to generate the needed fitness; smaller populations died out before they could do so. For all panels, the *environmental richness* $R = .005$.

Variable population size simulations. Figure 3 shows the results of simulations in which *population size N* varied across generations based on the population's fitness (mean *energy level E*). When the threshold T_N was low (Fig. 3A), relatively little fitness was required for the population to grow. If the population started with a small value of *N*, it began to evolve a high level of *controlled processing* due to limited competition. This in turn increased fitness, which increased the size of the population. However, as the population grew, the increased pressure due to competition led to the reemergence of automatic processing and diminution of control. Populations that started with a large *N* reached the same equilibrium, but the high levels of competition were already present and *controlled processing* did not become prevalent at any point.

With intermediate T_N , it was more difficult to sustain a large population (Fig. 3B). In this case, small and large starting populations alike reached an equilibrium characterized by a moderate degree of *controlled processing* (for small starting populations, this equilibrium was achieved only after *controlled processing C* peaked and then diminished).

When T_N was high, it was difficult for populations to achieve sufficient mean fitness to grow, or even maintain their size (Fig. 3C). Only the efficient resource consumption conferred by *controlled processing* could yield such a high mean fitness. As a consequence, small starting populations died out before they could evolve a sufficient prevalence of *controlled processing*. Larger populations, however, could sustain greater losses before dying out, providing enough time for *controlled processing* to evolve and go to fixation. Accordingly, larger populations evolved *controlled processing* quickly enough not only to halt the shrinking of *population size N*, but to increase *N* once they approached an equilibrium value near $C = 1$. Although the aforementioned simulations suggest that a large population of entirely automatic agents could not grow to such a size in the simulated environment, this does not preclude the invasion of that environment by one or more such populations. Thus, we see that when population size increases

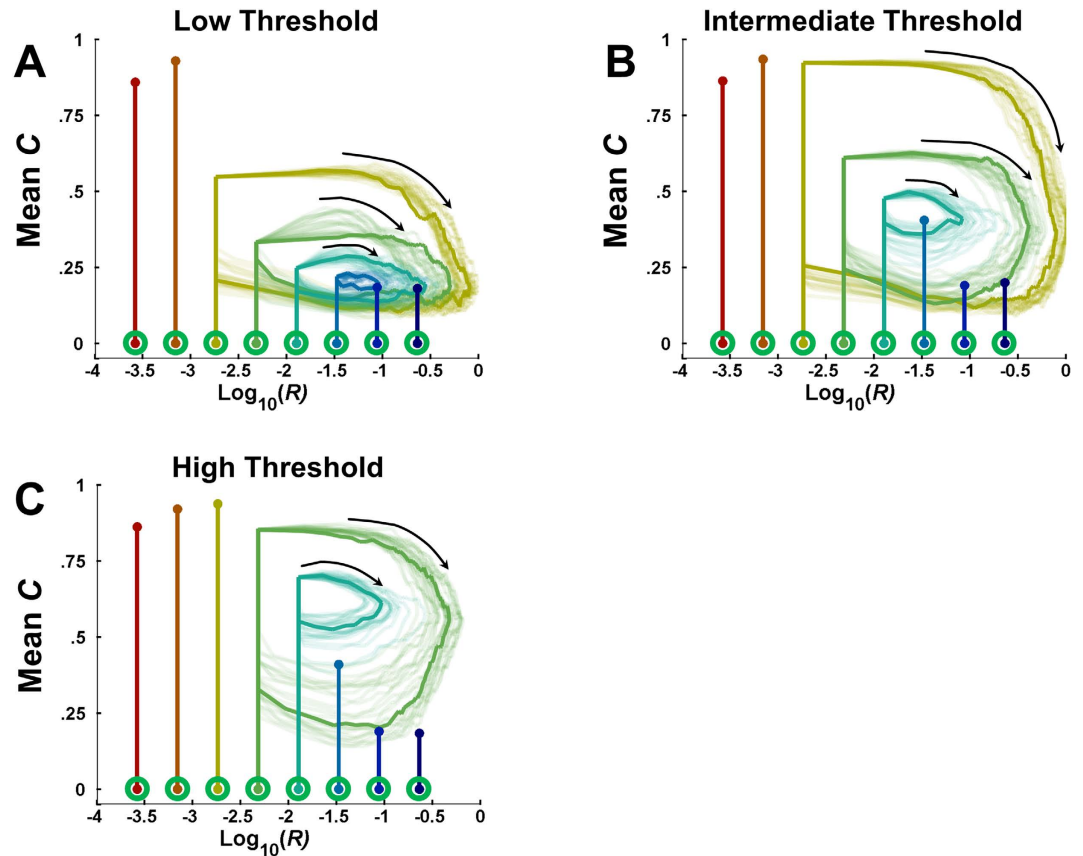


Figure 4. Variable richness simulations. The lines in each panel shows evolutionary trajectories, starting with *controlled processing* $C=0$ but with incrementally higher values of R_0 (x-axis). For those populations that exhibit limit cycles, many cycles are depicted and a single exemplar is shown with a darker line. Population size N was fixed at 10,000 for all simulations. (A) When the performance threshold T_R was small and several populations were able to modify *environmental richness* R . (B) A larger T_R reduced the number of populations able to modify R , but increased the range of the resulting limit cycles. (C) When T_R was large, only two populations achieved sufficient *mean energy levels* E and mean C to modify R .

with population fitness, a boom-bust pattern can occur in which controlled processing initially invades and grows the population, but then is overwhelmed by automatic processing as competition becomes too high.

Variable environmental richness simulations. Figure 4 shows the results of simulations in which R varied according to the combination of the population's mean *energy level* E and the prevalence of levels of *controlled processing* C . In lean environments where the initial richness R_0 was low (left side of each panel), even populations that evolved controlled processing lacked the fitness necessary to enrich the environment. Conversely, when R_0 was high (right side of panels), there was insufficient competitive advantage to controlled processing for it to evolve, and so there was little change to the environment.

For intermediate values of R_0 (middle sections of each panel), dynamics arose that were more interesting and complex. Initially, C increased, in turn increasing E . Once $E \cdot C$ was sufficiently large, R also began to increase. But as R grew, controlled processing became less beneficial, thus allowing automatic processing to proliferate and R to decline. The system returned to a leaner environment with its nominal richness R_0 , which again favored controlled processing, restarting the cycle. This effect was relatively robust to T_R , exhibiting limit cycles for intermediate values of R_0 at three different levels of T_R (Fig. 4A–C).

Discussion

We have presented a simulation-based evolutionary model in which agents used either of two types of processes to decide how to consume finite resources. Automatic processes were faster, providing an advantage when directly competing with other agents over a resource, but they followed a rigid consumption policy and thus fared poorly in leaner environments. In contrast, controlled processes were flexible, anticipating future needs and performing well in leaner environments. Nevertheless, when the advantages of controlled processing to the individual also had external consequences, such as increasing population size or resource abundance, these effects could diminish and even reverse the spread of

controlled processing through the population. We used two different scenarios to show that the success of controlled processing can undermine its own prevalence, leading to the re-invasion of automatic processing and/or limit cycles in the prevalence of processing style of an evolutionary timescale. These non-linear dynamics do not reflect simply a property of either automatic or controlled processing, but rather a complex interaction between the consequences of types of processing and the environment. Furthermore, these effects existed over a range of values for the parameters implemented in our simulations, and thus did not depend critically upon the conditions of any single set of initial conditions.

In one scenario, we assumed that higher mean fitness increased the size of the population, while lower mean fitness caused the population to shrink. When maintaining the population's size was relatively difficult (Fig. 3C), populations either died out or evolved sufficient levels of control to sustain themselves. When increasing the population's size was relatively easy, the excess fitness generated by controlled processing allowed the population to grow. However, the increase in competition at large population sizes led to the diminution (Fig. 3B) or demise (Fig. 3A) of control. Thus, while controlled processing may be favored to become prevalent, it may also be self-limiting due to its externalities and downstream consequences.

In a second scenario, controlled processing enriched the environment. These simulations yielded dynamics that were also non-monotonic, and somewhat more complex, reflecting the interaction between types of processing, competition, and resource availability. In this case, limits to the evolutionary success of controlled processing were due to free-loader-like effects: enrichment of the environment made it possible for agents using automatic processing to flourish and eventually outcompete those using controlled processing, at least for a time, by taking advantage of the greater resource availability provided by controlled agents (Fig. 4).

In our model, agents employing automatic processing used a simple, rigid policy: they completely consumed any resources they found. Moreover, automatic agents did not have access to stored goods. In additional simulations, we relaxed these assumptions, allowing agents using automatic processing to consume resources according to a parameter that evolved over generations in the same manner as the probability of controlled processing, C . This parameter dictated a “target energy level”: if the agent's current energy level was below this target level, the agent would consume any resources currently available (including those previously stored) until the current energy level attained the target level (in our main simulations, this target energy level was fixed to the maximum level, such that automatic processing always dictated full consumption). In these additional simulations, the automatic processing policy could adapt on an evolutionary time scale to best match the *environmental richness* R . Even under this framework, the rise and fall of controlled processing still occurred during variable population size simulations (see Supplementary Figure S2) and limit cycles still occurred in variable richness simulations (see Supplementary Figure S3). Thus, our main results are not unique to the assumption that automaticity involves immediate full consumption or that it cannot evolve, but instead are driven by the more general phenomenon of controlled processing adjusting to the environment more quickly than automatic processing. It is worth noting that such adaptations in automatic processing would only apply for that environment, and could not keep pace with environmental change that exceeded the pace of evolution (including those produced by the externalities of control). In fact, adverse consequences could ensue if the downstream effects of controlled processing outpace these forces^{21,44}, as suggested by existence of large limit cycles over in the variable richness simulations. Sufficient adaptation could also occur via cultural transmission, which is also cumulative but more rapid than genetic adaptation⁴⁵. Future work should investigate this idea more directly by examining the interaction between types of processing and the pace of both evolution and environmental change.

Taken together, the non-monotonicity of outcomes we observed under very different conditions and mechanisms supports the idea that flexibility at the individual level, while optimal for an individual in isolation, may not be optimal at the population level when it must compete with more rigid but more efficient strategies. These observations may have important scientific and practical implications, insofar as the balance between automatic and controlled processing influences decision making behavior in domains such as diet, drug use, savings, and the consumption of natural resources — behaviors that have consequences for both the individual but also the population and the environment.

These results make two significant contributions to the study of dual-process models of cognition. First, our framework establishes the utility of population models for studying the evolutionary trajectories of cognitive processes. Second, our simple case of evolutionary competition between controlled and automatic processing serves as a benchmark for future work that could consider more sophisticated forms of control. Such work might include control strategies that respond to, or even anticipate, the social and environmental effects of controlled processing, thereby potentially stabilizing the population at higher levels of controlled processing (and correspondingly higher mean levels of fitness).

The models we describe are relatively simple in other respects as well. More detailed models could consider environments with non-uniform spatial structure⁴⁶, in which agents could flexibly adapt to localized distributions of resources to produce “cultural topologies” of controlled vs. automatic processing. It may also be useful to explore cases in which controlled processing uses heuristics to speed processing when facing competition⁴⁷, anticipates and exploits the vulnerabilities of automatic processing (e.g., the role of marketing in promoting obesity), explicitly recognizes the value of prosocial behavior (e.g., cooperation⁴¹), is subject to distortions caused by social circumstances⁴⁸, or is constrained sociologically

(e.g., affects only others that are also using controlled processing). However, such strategies may carry other costs (e.g., energetic).

A further simplifying assumption of our simulations was that the controlled agents had access to the optimal consumption policy. Computing this optimal policy used a method that almost certainly exceeds the capabilities of real world agents. However, searching the space of possible actions need not be so exhaustive: computational architectures have been proposed that can dramatically limit search and still closely approximate optimal policies^{49,50}. Experiments have also shown that humans can effectively and unconsciously prune such a space with simple heuristics⁵¹.

It is worth noting that our simulations focused on a particular application of cognitive control: the problem of intertemporal choice (i.e., deciding between immediate and delayed consumption). There is a growing body of work suggesting that cognitive control plays an important role in such decisions^{14,15,25,26,32,52} and, though it is only one domain in which cognitive control operates^{2,3,16,53–55}, it may be a particularly important one. Intertemporal choice is fundamental to self-control: the capacity to behave in accord with long-term objectives despite temptations that have immediate benefits but long-term costs^{11,40,56,57}. Self-control, in turn, is at the heart of many social and environmental issues⁵⁸. Formal theory indicates that any rational account of self-control must posit at least two processes, one of which represents the agent's immediate interests and the other (the controlled process in our model) that serves its long-term interests⁴⁰. Thus, models that address phenomena having to do with self-control at the population level must be composed of agents built on at least a dual- (and possibly multi-) process model of decision making. The models we have described provide initial examples of this approach. More work, however, is needed to characterize the behavior of such systems in a fuller and more precise way, to incorporate more complex forms of control, and to validate this theoretical work in an empirical setting.

In summary, we introduced simulations exploring the evolutionary consequences of dual-system models of cognitive function, in which automatic and controlled processes competed for expression. We demonstrated that cognitive control, despite its capacity for flexible behavior and higher overall fitness, can evolve along a non-monotonic course as a result of its external influences on the population and/or the environment. That is, its initial spread can produce conditions that undermine its further evolution, and in some cases bring about its collapse. This type of model may be useful for exploring more sophisticated forms of controlled processing, and more complex interactions between individual decision making behavior and social or ecological factors. There are multiple domains in which different styles of human cognition can impact our wellbeing and that of our environment, and the computational modeling of large populations that embody these cognitive strategies is potentially applicable to a number of the problems posed by the very success of controlled processing in our own species.

References

- Cohen, J. D., Dunbar, K. & McClelland, J. L. On the control of automatic processes: a parallel distributed processing account of the Stroop effect. *Psychol. Rev.* **97**, 332–361 (1990).
- Miller, E. K. & Cohen, J. D. An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* **24**, 167–202 (2001).
- Posner, M. I. & Snyder, C. R. R. *Attention and cognitive control in Information processing and cognition: the Loyola symposium* (ed. Solso R. L.) 55–85 (L Erlbaum Associates, Hillsdale, N. J., 1975).
- Schneider, W. & Shiffrin, R. M. Controlled and automatic human information processing: I. detection, search, and attention. *Psychol. Rev.* **84**, 1–66 (1977).
- Shiffrin, R. M. & Schneider, W. Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychol. Rev.* **84**, 127–190 (1977).
- Barrett, H. C. & Kurzban, R. Modularity in cognition: framing the debate. *Psychol. Rev.* **113**, 628–647 (2006).
- Evans, J. S. B. T. In two minds: dual-process accounts of reasoning. *Trends Cogn. Sci.* **7**, 454–459 (2003).
- Evans, J. S. B. T. & Stanovich, K. E. Dual-process theories of higher cognition: advancing the debate. *Perspect. Psychol. Sci.* **8**, 223–241 (2013).
- Stanovich, K. E. & West, R. F. Individual differences in reasoning: implications for the rationality debate? *Behav. Brain Sci.* **23**, 645–665 (2000).
- Tversky, A. & Kahneman, D. Extensional versus intuitive reasoning: the conjunction fallacy in probability judgment. *Psychol. Rev.* **90**, 293–315 (1983).
- Fudenberg, D. & Levine, D. K. A dual-self model of impulse control. *Am. Econ. Rev.* **96**, 1449–1476 (2006).
- Kahneman, D. Maps of bounded rationality: psychology for behavioral economics. *Am. Econ. Rev.* **93**, 1449–1475 (2003).
- Thaler, R. H. & Shefrin, H. M. An economic theory of self-control. *J. Polit. Econ.* **89**, 392–406 (1981).
- McClure, S. M., Laibson, D. I., Loewenstein, G. & Cohen, J. D. Separate neural systems value immediate and delayed monetary rewards. *Science* **306**, 503–507 (2004).
- Hare, T. A., Camerer, C. F. & Rangel, A. Self-control in decision-making involves modulation of the vmPFC valuation system. *Science* **324**, 646–648 (2009).
- Ariely, D. *Predictably irrational* (HarperCollins LLC, New York, 2008).
- Loewenstein, G. Out of control: visceral influences on behavior. *Organ. Behav. Hum. Decis. Process.* **65**, 272–292 (1996).
- Bialek, W., Nemenman, I. & Tishby, N. Predictability, complexity and learning. *Neural Comput.* **13**, 2409–2463 (2001).
- Ernst, M. O. & Banks, M.S. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* **415**, 429–433 (2002).
- Kording, K. P. & Wolpert, D. M. Bayesian integration in sensorimotor learning. *Nature* **393**, 266–268 (2004).
- Cohen, J. D. The vulcanization of the human brain: a neural perspective on interactions between cognition and emotion. *J. Econ. Perspect.* **19**, 3–24 (2005).
- Cosmides, L. & Tooby, J. Beyond intuition and instinct blindness: toward an evolutionarily rigorous cognitive science. *Cognition* **50**, 41–77 (1994).
- Kahneman, D. *Thinking, fast and slow* (Farrar, Straus and Giroux, New York, 2011).
- Gould, S. J. & Lewontin, R. C. The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist programme. *Proc. R. Soc. Lond. B Biol. Sci.* **205**, 581–598 (1979).

25. Figner, B. *et al.* Lateral prefrontal cortex and self-control in intertemporal choice. *Nat. Neurosci.* **13**, 538–539 (2010).
26. Van den Bos, W. & McClure, S. M. Towards a general model of temporal discounting. *J. Exp. Anal. Behav.* **99**, 58–73 (2013).
27. Hinson, M., Jameson, T. L. & Whitney, P. Impulsive decision making and working memory. *J. Exp. Psychol. Learn. Mem. Cogn.* **29**, 298–306 (2003).
28. Ward, A. & Mann, T. Don't mind if I do: disinhibited eating under cognitive load. *J. Pers. Soc. Psychol.* **78**, 753–763 (2000).
29. Ainslie, G. Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychol. Bull.* **82**, 463–496 (1975).
30. Laibson, D. Golden eggs and hyperbolic discounting. *Q. J. Econ.* **112**, 443–478 (1997).
31. Thaler, R. Some empirical evidence on dynamic inconsistency. *Econ. Lett.* **8**, 201–207 (1981).
32. Duncan, J. & Owen, A. M. Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends Neurosci.* **23**, 475–483 (2000).
33. Preuss, T. M. Do rats have prefrontal cortex? The Rose-Woolsey-Akert program reconsidered. *J. Cogn. Neurosci.* **7**, 1–24 (1995).
34. Shallice, T. & Burgess, P. The domain of supervisory processes and temporal organization of behavior. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **351**, 1405–1412 (1996).
35. Uylings, H. B. M. & van Eden, J. P. C. Qualitative and quantitative comparison of the prefrontal cortex in rat and in primates, including humans. *Prog. Brain Res.* **85**, 31–62 (1990).
36. Grether, D. M. & Plott, C. R. Economic theory of choice and the preference reversal phenomenon. *Am. Econ. Rev.* **69**, 623–638 (1979).
37. Tversky, A., Slovic, P. & Kahneman, D. The causes of preference reversal. *Am. Econ. Rev.* **80**, 204–217 (1990).
38. O'Donoghue, T. & Rabin, M. Doing it now or later. *Am. Econ. Rev.* **89**, 103–124 (1999).
39. Strotz, R. H. Myopia and inconsistency in dynamic utility maximization. *Rev. Econ. Stud.* **23**, 165–180 (1955).
40. Gul, F. & Pesendorfer, W. Temptation and self-control. *Econometrica* **69**, 1403–1435 (2001).
41. Rand, D. G., Greene, J. D. & Nowak, M. A. Spontaneous giving and calculated greed. *Nature* **489**, 427–430 (2012).
42. Kahneman, D. & Treisman, A. Changing views of attention and automaticity in *Varieties of Attention* (eds Parasuraman, R. & Davies, D. R.) 29–61 (Academic Press, Orlando, 1984).
43. Ewens, W. J. *Mathematical Population Genetics I: Theoretical Introduction* (Springer, New York, 2004).
44. Baronchelli, A., Chater, N., Christiansen, M. H. & Pastor-Satorras, R. Evolution in a changing environment. *PLoS One* **8**, e52742 (2013).
45. Tomasello, M. Human culture in evolutionary perspective in *Advances in Culture and Psychology* (eds Gelfand, M., Chiu, C., Hong, Y.) 5–51 (Oxford University Press, Oxford, 2011).
46. Moreira, J. A., Pacheco, J. M. & Santos, F. C. Evolution of collective action in adaptive social structures. *Sci. Rep.* **3**, 1521 (2013).
47. Gigerenzer, G., Todd, P.M. & ABC Research Group. *Simple heuristics that make us smart* (Oxford University Press, Inc., New York, 1999).
48. Johnson, D. D. P. & Fowler, J. H. The evolution of overconfidence. *Nature* **477**, 317–320 (2011).
49. Kearns, M., Mansour, Y., Ng, A. Y. A sparse sampling algorithm for near-optimal planning in large Markov decision processes. *Mach. Learn.* **49**, 193–208 (2002).
50. Sutton, R. S. Integrated architectures for learning, planning, and reacting based on approximate dynamic programming in *Proceedings of the seventh international conference on machine learning* 216–224 (1990).
51. Huys, Q. J. M. *et al.* Bonsai trees in your head: how the Pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Comput. Biol.* **8**: e1002410 (2012).
52. Metcalfe, J. & Mischel, W. A hot/cool-system analysis of delay of gratification: dynamics of willpower. *Psychol. Rev.* **106**: 3–19 (1999).
53. Evans, J. S. B. T. Dual-processing accounts of reasoning, judgment, and social cognition. *Annu. Rev. Psychol.* **59**, 255–278 (2008).
54. Hofmann, W., Schmeichel, B. J. & Baddeley, A. D. Executive functions and self-regulation. *Trends Cogn. Sci.* **16**, 174–180 (2012).
55. Posner, M. I. & Petersen, S. E. The attention system of the human brain. *Annu. Rev. Neurosci.* **13**, 25–42 (1990).
56. Mischel, W., Shoda, Y. & Rodriguez, M. L. Delay of gratification in children. *Science* **244**, 933–938 (1989).
57. Baumeister, R. F. & Heatherton, T. F. Self-regulation failure: an overview. *Psychol. Inq.* **7**, 1–15 (1996).
58. Santos, F. C. & Pacheco, J. M. Risk of collective failure provides an escape from the tragedy of the commons. *Proc. Natl. Acad. Sci. U. S. A.* **108**, 10421–10425 (2011).

Acknowledgments

The authors thank S. Levin, A. Shenhav and J.M. White for their suggestions regarding this manuscript, and E. Maskin, S. Levin and A. de Froment for early discussions of this project. This project/publication was made possible through the support of a grant from the John Templeton Foundation. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

Author Contributions

D.T., D.G.R., E.A.L. and J.D.C. designed the agent-based simulations. E.A.L. designed the optimal consumption algorithm. D.T. programmed the simulations, collected and analyzed the data. D.T., D.G.R., E.A.L. and J.D.C. prepared the manuscript.

Additional Information

Supplementary information accompanies this paper at <http://www.nature.com/srep>

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Tomlin, D. *et al.* The evolution and devolution of cognitive control: The costs of deliberation in a competitive world. *Sci. Rep.* **5**, 11002; doi: 10.1038/srep11002 (2015).



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>