

# Choosing Fitness-Enhancing Innovations Can Be Detrimental under Fluctuating Environments

Julian Z. Xue<sup>1\*</sup>, Andre Costopoulos<sup>2</sup>, Frederic Guichard<sup>1</sup>

<sup>1</sup> Department of Biology, McGill University, Montreal, Canada, <sup>2</sup> Department of Anthropology, McGill University, Montreal, Canada

## Abstract

The ability to predict the consequences of one's behavior in a particular environment is a mechanism for adaptation. In the absence of any cost to this activity, we might expect agents to choose behaviors that maximize their fitness, an example of directed innovation. This is in contrast to blind mutation, where the probability of becoming a new genotype is independent of the fitness of the new genotypes. Here, we show that under environments punctuated by rapid reversals, a system with both genetic and cultural inheritance should not always maximize fitness through directed innovation. This is because populations highly accurate at selecting the fittest innovations tend to over-fit the environment during its stable phase, to the point that a rapid environmental reversal can cause extinction. A less accurate population, on the other hand, can track long term trends in environmental change, keeping closer to the time-average of the environment. We use both analytical and agent-based models to explore when this mechanism is expected to occur.

**Citation:** Xue JZ, Costopoulos A, Guichard F (2011) Choosing Fitness-Enhancing Innovations Can Be Detrimental under Fluctuating Environments. PLoS ONE 6(11): e26770. doi:10.1371/journal.pone.0026770

**Editor:** James A.R. Marshall, University of Sheffield, United Kingdom

**Received:** June 29, 2011; **Accepted:** October 3, 2011; **Published:** November 17, 2011

**Copyright:** © 2011 Xue et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Funding:** This work was partially funded by a Fonds de recherche sur la société et la culture Team grant on hominid dispersals, the European Science Foundation's BOREAS Histories From the North program, Social Sciences and Humanities Research Council of Canada, and Canadian Institutes of Health Research. No additional external funding was received for this study. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

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

\* E-mail: zhunping.xue@mail.mcgill.ca

## Introduction

One of the distinguishing features of biological evolution is the production of adaptation through natural selection on blind variation [1]. A vast amount of theory and empirical work has been dedicated to this topic, stemming from the earliest works of Fisher [2] and other classical population geneticists (e.g. [3,4]). However, this mechanism is not the only method by which adaptation can be produced.

Population genetics is fairly sparse on the subject of directed adaptation, or the overproduction of new phenotypes that adapt the organism well to the current environment. This is, of course, because genetic mutation is blind. However, the behavioral innovations of some animals are not necessarily blind. Some animals can predict the consequences of their actions within their environment, which increases their probability of choosing a behavior or innovation among a large pool of possibilities that fit them well to the current environment. These behaviors can then be transmitted via cultural inheritance to the next generation. Thus they can adapt behaviorally in a directed way.

Intuitively, we can believe that such directedness promotes the maintenance of high fitness. In fact, the scholarly work on the capacity for predicting the consequences of our actions is littered with optimistic statements, going all the way back to Malthus, who spoke of:

"...that distinctive superiority in [human] reasoning faculties, which enables [them] to calculate distant consequences" [5].

The intuition that being able to predict the consequences of one's actions is good for fitness remain today in the works of most

authors on this topic (e.g. [6,7] and comments in [7]). For example, D'Argebeau in a comment to [7] says:

"S&C [Suddendorf and Corballis] argues that the primary fitness function of mental time travel is to enhance biological fitness in the future: Mentally simulating various versions of the future, and their respective consequences, enables one to act flexibly in the present to increase one's future survival chances. We completely agree..."

The reasoning seems to be that agents with the ability to predict the consequences of their action in the environment can choose the best possible action out of the pool of all possible actions, and thereby increase its adaptedness, and therefore its fitness. However, these predictions ignore the effects of environmental fluctuations, including sudden changes in fitness optima, which often have counterintuitive consequences for evolution [8,9]. One would like a more rigorous argument of when the directedness of innovations is advantageous, as well as quantitative predictions of just how advantageous [10]. After all, it is reasonable to assume that the capacity for directed innovation increases with brain size, which is expensive in terms of energy [11], obstetric difficulties [12], and a long childhood [13]. A clearer theoretical understanding would allow us to predict trade-offs involving directed innovation, as well as the situations in which directed innovation can be more, or less, important.

Here, we derive a scenario where maximizing one's accuracy in choosing the fittest innovation in the current environment can be detrimental. The systems we study are those with dual inheritance [14], where the parent is able to pass on both genes as well as learned behavior. Our approach is based on the intuition that

agents that are too good at choosing fitness-enhancing innovations can overfit environmental change. This possibility has already occurred to Darwin, who wrote that

“If animals became adapted to every minute change, they would not be fitted to the slow great changes really in progress.” [15].

Thus, agents who are capable of assessing the consequences of their own actions in the immediate environment, but who are incapable of forecasting the changes of the environment, might be thrown off by environmental noise rather than adapting themselves to long term trends or time-averages. We more specifically show how the occurrence of sudden changes in the environment can limit the degree of directed innovation. We first provide an intuitive description of this scenario. We then present a formal analytical treatment, and finally we report results from an agent-based simulation applied to a historical example to which our scenario might apply.

## Models and Results

### Toy model

Let us begin with an oversimplified toy model in order to formalize this intuition. Consider an environment, characterized by a single parameter  $E$  that varies around a certain mean. In fact, let us consider a very particular type of environment, where  $E$  slowly increases over a long period of time, then quickly reverses and crashes (see Figure 1).

Let's say there are two agents exposed to  $E$  and there is a pool of actions from which each of the agents can choose. Different actions alter their fitness in response to  $E$ . However, the pair of agents differ in their ability to assess the consequences of any particular action in the environment. We call the one who can assess better the “accurate” agent and the one who assess less well the “inaccurate” agent. Note that this is not yet a model of evolution, rather of competition between two agents who do not die. The competition is based on how well the agents learn from the environment.

We model the capacity for directed innovation as the average reduction of the distance between individual fitness and the fitness optimum following each innovation. This is intuitively reasonable because the agents who can assess better can change its behavior to be more in tune with changes in the environment, and hence produce better adaptations. The inaccurate agent, on the other hand, cannot adapt its behavior to environmental change so well.


Consider the case where the inaccurate agent is so poor that, on average, its innovations cannot even track the environmental changes during the phase of slow increase in  $E$ , and remain near the long term environmental mean despite the slow environmental change. Thus, for the period of slow environmental change, the accurate agent will do better. However, when the environment rapidly reverses, the accurate agent might in fact do worse. It might do worse because it fits the previous change so well that the rapid reversal meant the environment moved further away from its optimum than the optimum of the inaccurate agent, which had stayed near the environmental mean.

We can reason out the conditions allowing inaccurate types to maintain a higher long-term fitness than accurate types. First of all, if the accurate agent is able to perfectly adapt to the phase of rapid environmental change, then it will always have a higher fitness than the inaccurate agent. The perfect strategy for agents is to perfectly track the environment. However, this strategy is impossible if environment changes so rapidly in its rapid phase as to escape the capacity of both types of agents to adapt.

Secondly, the environment must possess a strong central tendency and be varying about some longer-term mean, or trend, since if its rapid phase is in the *same* direction as the slow phase, then the accurate agent will be superior. Finally, the fitness function must be non-linear in the sense that the cost to the accurate agent after the rapid environmental reversal must outweigh the fitness advantage it has accumulated over the period of slow environmental change.

The first assumption, that perfect adaptation is impossible, seems to be a reasonable one in view of how the environment changes at all scales of time. It is impossible for organisms to adjust themselves to every minute change that occurs with great rapidity. The second one also seems reasonable, since most environmental conditions do vary about a certain mean, although whether they change slowly in one direction and then quickly in another depends on the scenario at hand, and we will give an example in the final section.

The third assumption is more subtle. Superficially, it already seems quite reasonable, since most fitness curves are convex near the optimum [16–19]. The usual theoretical curve used here is a gaussian curve, so both theoretically as well as empirically the fitness drops more precipitously as the agent's phenotype moves away from the environmental optimum. Thus, a short moment of severe maladaptation can strongly outweigh long periods of relative better adaptation. On the other hand, a more subtle



**Figure 1. An artificial environment used by the agent based model.**  $E(t)$  represents the environmental value at time  $t$ . doi:10.1371/journal.pone.0026770.g001

theoretical point is that the curvature of the fitness curve is not necessary, although a strong curvature makes the less accurate type much more likely to be more fit. This is due to the well-known principle that natural selection maximizes the geometric mean fitness over time, not the arithmetic mean [20–23]. Thus, while maximizing fitness, natural selection also minimizes variance in fitness over time, and since the worst fitness of the accurate agent is worse than the worst fitness of the inaccurate agent, the accurate agent accrues more variance in its fitness, lowering its geometric mean fitness over time.

We next construct an analytical model to provide a formal theory of maladaptive directed innovation under sudden environmental changes.

### Analytical model

The above section describes a toy model in which maximizing the fitness of each decision may not be an optimal strategy when the environment changes slowly in one direction for long periods and then rapidly reverses. This can be illustrated with a numerical example. Consider a simple cyclical environment with cycle 4 that looks like the following:  $E(t) = \{0, 1, 2, -1, 0, 1, 2, -1, \dots\}$ , where  $t$  indicates the timestep.

Consider two agents,  $A$  and  $B$ , where  $A$  is more accurate than  $B$ . Due to greater accuracy,  $A$  can adapt to the environment more quickly than  $B$  by choosing better actions. Let us say, for example,  $A$  can change its optimal phenotype by a maximum of  $x_A = 1$  per time step in the direction that will take its phenotype closer to the environmental optimum.  $B$ , on the other hand, is incapable of choosing so well and  $x_B = 0.5$ . Let both agents be perfectly adapted to the environment initially, and let  $\phi(t)$  be the phenotype of the agent at time  $t$ . The phenotype is measured as the environment to which the organism is best adapted. We thus set  $E(0) = \phi_A(0) = \phi_B(0) = 0$ .

We give the first few entries of the sequence of  $E(t)$ ,  $\phi_A(t)$ , and  $\phi_B(t)$  for clarity:

$$E(t) = \{0, 1, 2, -1, 0, 1, \dots\}$$

$$\phi_A(t) = \{0, 1, 2, 1, 0, 1, \dots\}$$

$$\phi_B(t) = \{0, 0.5, 1, 0.5, 0, 0.5, \dots\}$$

We introduce the fitness curve

$$\lambda = \frac{1}{|E(t) - \phi(t)|^\gamma + 1} \tag{1}$$

where  $\gamma$  is a tunable parameter. This fitness curve looks a great deal like a Gaussian curve, but is easier to compute with (Figure S1, see Information S1).  $\gamma$  controls the rate at which the curve falls off as  $|E(t) - \phi(t)|$  increases; a greater  $\gamma$  means a sharper threshold of  $|E(t) - \phi(t)|$  at which fitness falls off precipitously.

It turns out that  $B$  is more fit than  $A$  for  $\gamma > 3.238\dots$ . In fact, the existence of some critical  $\gamma_c$  is true in general; for similar environments and fitness curves, then an agent  $B$  with less foresight than agent  $A$  will be more fit than agent  $A$  if  $\gamma$  is large than some  $\gamma_c$  (see Information S1). Moreover, we can show that for any such environment (such as the one in Figure S2), agents behave cyclically and there is an intermediate value of  $x$  which is locally optimal for any  $\gamma$ . (See Figure S3, Information S1).

### Agent-based model

Our model of directed innovation differs from models of blind mutation. In blind models, mutations themselves do not take the agent closer to the environmental mean. Any movement towards the environmental mean is due to the differential effects of natural selection. In the above analytical model, however, the directed movement is purely due to the capacity of the agent to choose, on average, an innovation that takes it closer to the environment. The previous model is not a model of evolution. We did not include variance in innovations and there was no reproduction or selection.

The analytical model shows that agents with less directedness can be more fit than agents with more directedness. In order to make the model evolutionary, we extend the above analytical model with a dual inheritance agent-based model, where the agent's behavior begins as being learned (culturally inherited) from its parent, then changes according to its genetically determined capacity to learn from the environment, which is genetically inherited from the parent. This allows us to evaluate the likelihood of mutating off the local optimum, severity of conditions required on  $\gamma$  (if any), robustness to drift effects, limited population sizes, and the applicability of this theoretical scenario to the real world. We confirm that the intermediate accuracy under the conditions described by the analytical model is a robust prediction under these scenarios.

In the agent-based model, as in the analytical model, individual agents have a phenotypic value ( $\phi$ ) that represents their adaptation to an environment. The environment is uniform with one dimension of variability described by  $E(t)$ . At each time step, behavioral innovations become available. Agents can adopt or refuse these innovations and as a result move their phenotype closer or farther away from the environmental optimum, increasing or decreasing their fitness. Agents have a second trait ( $x$ ) that represents their probability of knowing how behavioral innovations will affect them. Accurate agents (higher  $x$ ) are better at determining whether any behavioral innovation will move their phenotypic value closer to the current environmental value. Inaccurate agents have lower  $x$ . Agents adopt the behavioral innovation if and only if they can determine that it is beneficial. The probability of adoption is  $xP_b$ , where  $P_b$  is the probability that a new behavior is beneficial. Agents reject the behavioral innovation if they are unable to determine whether it is beneficial, or if they determine it to be harmful. The assumption that new behaviors are always rejected if agents cannot foresee their consequences reduces the noise in simulation results. Otherwise inaccurate agents would have a great range of fitnesses (depending on what behaviors they adopt by chance), and these fitnesses would be independent of their foresight. This essentially turns the model into one of blind mutation, which we would like to avoid.

Consider  $N$  agents over the time period  $[0, T]$ , and the environment changes as  $E(t)$ . Each agent  $i$  is defined by two parameters,  $\{\phi_i, x_i\}$ , where  $\phi_i$  is the environment that  $i$  is adapted to, and  $x_i$  is the probability that agent  $i$  can assess the consequences of adopting a behavioral innovation, that is, whether the new behavior will move the agent closer or farther from the environmental optimum. At each timestep, there is a probability,  $f_{innov}$ , that a new behavior will be available for adoption. This behavior will increase or decrease  $\phi_i$  by  $s_{innov}$  with equal probability (so  $P_b = 0.5$ ) if adopted by any given agent.  $s_{innov}$  thus gives an upper bound to the maximum rate of adaptation to the current environment.

At each timestep,  $N$  agents are chosen with replacement for update. During each update, agent  $i$  survives with the probability

$$\lambda = \frac{1}{|E(t) - \phi(t)|^\gamma + 1} \quad (2)$$

We used survival instead of reproductive fitness to increase simulation speed, but preserve the same qualitative behavior of the fitness curves [24]. The dynamics of the model are robust to implementation scheme; for example, a random agent can be chosen to replicate with this fitness, pushing out another random agent.


In our implementation, if agent  $i$  does not survive, a random agent  $j$  (not including  $i$ ) is chosen from the population to reproduce. Its progeny then replaces  $i$ . Thus, a new agent  $i^*$ , a progeny of agent  $j$ , replaces the agent  $i$  that has just died. The new agent  $i^*$  is very similar to  $j$  with small modifications. Specifically,  $\phi_{i^*} = \phi_j$ , which is an act of cultural transmission (the passage of a learned behavior from parent to progeny), and  $x_{i^*} = x_j + \theta$ , an act of genetic transmission with mutation, where  $\theta \in [-\sigma, \sigma]$ , a uniform distribution.  $x_{i^*}$  is a probability that is kept bounded between 0 and 1. We do not allow mutations in the culturally learned behavior because we wish for all movement of the phenotype towards the environmental mean to be a result of directed innovations by choice of the agents, rather than blind chance and selection.

If agent  $i$  survives and a new behavior is available, it has  $x_i$  chance of accurately assessing whether the new behavior will take it closer to the environmental optimum. If it assesses successfully, and the new behavior moves the agent closer to the environment, then it adopts the behavior and modifies its  $\phi_i$ . Else it does not adopt the new behavior. This is how we model directed innovations.

We tested the model with a variety of environments. We report the results with an artificial environment (Figure 1) and one based on the Vostok ice core data [25,26] (see next Section). In order to make the different environments comparable, we normalize them so the the minimum environmental value is 0 and the maximum environmental value is 1. We examine a range of  $\gamma$ 's and we note the minimum  $\gamma$  value for which the agents settle into an intermediate, rather than a maximal ability to foresee the consequences of new behaviors. For all simulations,  $N = 10000, f_{innov} = 0.1, s_{innov} = 0.005, \sigma = 0.01$ . For the artificial environment presented in Figure 1, predicted  $x$  values are independent of initial conditions for all  $\gamma \geq 6$ . The results for  $\gamma = 5, \gamma = 6, \gamma = 10$  are presented in Figure 2. As expected from the analytical model, for large enough  $\gamma$  ( $\gamma > 6$ ) agents do not evolve to maximum possible foresight. Each sudden environmental reversal will favour anew the inaccurate agents.

This result depends on the assumption that even maximum foresight is not sufficient to perfectly track the environment during the sudden reversal. If perfect tracking was possible, the agents that could perfectly track would necessarily be the most fit. We argued in our Introduction that the inability to perfectly track the environment during sudden environmental changes is a reasonable assumption, and we will return to this assumption in the Discussion.

The main differences between the agent-based and the analytical models presented here are the dissociation between  $x$  and  $\phi$ , and limited population size introduced in the agent-based model. If the environment slowly rises in one direction for too long, or if  $\gamma$  is too small, then all agents with low  $x$  die before the onset of sudden environmental change, in which case high  $x$  dominates, or the population goes extinct (which is prevented in this model). The minimum  $\gamma$  values needed to generate



**Figure 2. Intermediate adaptive rate ( $x$ ) is optimal for agents in the environment of Figure 1.** Brightness denotes density of agents. From top to bottom:  $\gamma = 5, \gamma = 6, \gamma = 10$ . All other parameters are identical (see text for more details). Initial conditions:  $E(0) = \phi_i(0) = 0, x_i \in [0, 1]$ . The results are independent of initial conditions. Note how the larger the gamma, the stronger the movement to intermediate, rather than maximal, accuracy.  
doi:10.1371/journal.pone.0026770.g002

intermediate optimal  $x$  from the model means that we need a fitness curve that drops off quickly when the distance between the phenotype and the environmental optimum reaches a threshold. This way, inaccurate agents can survive until the catastrophic event (since they do not cross that threshold), but the catastrophic event is sufficient to kill off most of the accurate agents (since they do cross that threshold). Thus, although the analytical model predicts a locally optimal  $x$  for all  $\gamma$ , a minimum  $\gamma$  is required in the agent-based model to observe a local optimum.

### Agent-based modeling using the Vostok ice core data

The artificial environment we use in the previous section can be argued to be too artificial; the rise and the fall are both clean of noise and the sudden reversal might be too sudden and too quick to be realistic. To address this issue, we insert a more realistic timeseries in the form of the Vostok ice core data [25,26]. This data indicates that the temperature environment over the last 400 000 years has been characterized by rapid deglaciation events, interspersed across periods of slower cooling and cold plateaus. We test whether this could result in an intermediate, locally optimal adaptive rate. We replace our artificial environment with the Vostok ice core data (Figure 3), where each timestep counts for a


single year. All other details remain identical to the previous model. Even  $\gamma=5$  in this model favored agents of moderate or even low ability to foresee the consequences of adopting different behaviors (Figure 4). As expected, the effect increases with increasing  $\gamma$ .

The first significant period of change is a rapid glaciation 400kya that causes inaccurate agents to have a fitness advantage. Four rapid deglaciation events 350kya, 250kya, 150kya, and a last one in the past 20ky, all allow low inaccurate agents to persist for long periods. In other words, there is an advantage to reducing the overall ability of the agents to discriminate between beneficial and harmful innovations under these circumstances. We believe that a large range of environmental variables display this pattern of slow movement in one direction interspersed with catastrophic reversals, which potentiates the type of scenario described in this paper. Natural disasters, for example, might happen quickly, followed by long and slow periods of recovery.


**Discussion**

In this paper, we show a somewhat counterintuitive point that agents which can best assess the consequences of their actions in the current environment can nevertheless be at a fitness disadvantage relative to less accurate competitors. We reason that the ability of agents to predict the consequences of their actions in a given environment acts like a greedy algorithm; it constantly tries to optimize for the moment. In the cases reported here, the greedy algorithm fails because the accurate agent over-optimizes for the immediate present, and so fails to track the long term average of the environment.

Several veins of earlier theoretical work shed light on these results. The first is the work on optimal rates of mutation (e.g. [27–30]). Another related line is natural selection and mutation rates in changing environments (e.g. [21,31,32]). The first body of work established that in the majority of cases, a low or zero mutation rate is preferable since the vast majority of mutations are detrimental. The fact that mutation rates are not zero is likely due to the physiological cost of increasing copying fidelity ([30]), especially in sexual beings where the hitchhiking of alleles that increases mutation rate to beneficial alleles is broken by recombination. The second line of work established the manner in which environments must change in order to maintain genetic



**Figure 3. The Vostok ice core data.** Long periods of cooling are punctuated by rapid deglaciation. doi:10.1371/journal.pone.0026770.g003



**Figure 4. Agent discrimination ability using the Vostok ice core data.** From top to bottom:  $\gamma=5, \gamma=6, \gamma=10$ . Notice how discriminatory ability rises during periods of slow change, but falls each time there is a rapid reversal. doi:10.1371/journal.pone.0026770.g004

variability (e.g. [32]). These results, however, are primarily focused on blind mutations. Directed innovations, where a larger proportion of change might be beneficial, require a different sort of analysis.

The closest work focuses on the evolution of individual learning (e.g. [8,9,14,33]). The evolution of learning from the environment is strongly related to this work, since this type of learning (as opposed to social learning from other agents) can be seen as the act of choosing new behaviors with feedback from the environment. Most of this work, however, has focused on the evolution of the capacity for culture. A recent result that studies learning in variable environments, for example, showed that in a stable environment punctuated by sharp bouts of change, agents stop learning from the environment during the stable periods and instead opt to choose to learn from each other. When the environment does change, the agents can no longer learn from the environment and may go extinct [9]. Our work instead focuses purely on the dynamics of the evolution of learning from the environment and on the optimal rates of such learning. Previous work assumed that better capacity for learning, if costless in other fitness terms such as time and energy, would likely lead to improvement of agent fitness, although our results have been foreseen by [34], who realized that changes in environmental conditions can lead to a reduction of learning.

Our model applies to systems of dual inheritance, where the learning done by the parent generation is passed on to the progeny through cultural transmission, and the capacity to learn further from the environment is passed on through genetic transmission. The model also contains three key simplifying assumptions that allow for analytical tractability and increase the generality of our results. First is that agents cannot fully adapt to the environment using directed innovation when environmental change is at its quickest. This, we argue, is reasonable. The environment fluctuates on every timescale; even organisms with learning cannot adapt to every change. In the absence of prior adaptation, the maximum rate of environmental change during a volcanic eruption, for example, must be very quick indeed and no cognitive capacity is likely to cope with it in an optimal way. A recent review of the evolution literature [35] falls in very neatly with this assumption, that the environment constantly exerts considerable selection pressure on organisms, but with many changes in direction and with longer trends of change that proceed at a much slower pace.

The second assumption is that agents are incapable of forecasting the environment. Forecasting denotes agents' knowledge of the future environment. If there are agents with good environmental forecasting abilities who can know the environment many timesteps in advance, they can preadapt and nullify these theoretical predictions. Future explorations will allow us to understand the sensitivity of this model to forecasting. However, we expect our results to be robust to some amounts of forecasting, which would be consistent with overall difficulty of predicting environmental fluctuations.

The third assumption of the model is that agents are incapable of remembering or reusing behavioral innovations they adopted in the past. This assumption is implicit in the fact that the agents' directed rates of adaptation are the same in any direction, whether or not they previously explored that parameter range. This assumption is realistic for many animal species using only limited forms of social learning, so behavioral innovations gained by ancestors are forgotten as quickly as they become unfit. The assumption holds when there is no repertoire of past behavior from which the current generation can draw to deal with environmental conditions that occurred in the past.

We make the important caveat that we do not believe that the evolution of directed innovation, particularly in the vastly more complex system of humans, in fact occurred in the way we

describe it here. Too many parameters, including the three assumptions stated above: the maximum rate of directed innovation, the existence of memory and forecasting abilities, remain unknown. We have, however, stated a theoretical possibility with clear and general implications. It can potentially provide insights into the cultural evolution of many animal species. Future work will provide more quantitative predictions that can be directly tested for particular species, cultural systems and environmental regimes.

## Supporting Information

**Figure S1 An example fitness curve for  $\gamma = 4$ ,  $E = 0$ , and  $\phi$  from  $-5$  to  $5$ .** It is very similar to the normal distribution curve but is easier to compute.

(TIFF)

**Figure S2 The environment used in the analytical model.** The environment increases in discrete steps of 1 each timestep for  $c$  timesteps, then drops to 0 in a single time step.

(TIFF)

**Figure S3 Agent behavior in this environment over the first cycle.** Dashed lines and diamonds indicate the phenotype of the agent, solid lines and circles the environment. The particular  $x$  (maximum rate of phenotypic change) used here is  $x = 0.5$ , but changing  $x$  will only change the amplitude of the cycle, not its phase or frequency.

(TIFF)

**Information S1.**

(PDF)

## Acknowledgments

We would like to thank the following colleagues for useful comments and discussion during the writing of this paper: Colin Chapman, Pete Richerson, Mark Lake, Axel Vandenberg, Luke Premo, Colin Wren. We would also like to thank two anonymous reviewers for their highly constructive feedback.

## Author Contributions

Conceived and designed the experiments: JX AC. Performed the experiments: JX. Analyzed the data: JX. Wrote the paper: JX AC FG.

## References

- Darwin C (1958) *The origin of species*. Hayes Barton Press.
- Fisher S (1999) *The genetical theory of natural selection: a complete variorum edition*. Oxford University Press, USA.
- Haldane J (1990) *The causes of evolution*. Princeton University Press.
- Wright S (1932) The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In: *Proc of the 6th International Congress of Genetics*. volume 1 356–366.
- Malthus T (1872) *An essay on the principle of population: or, A view of its past and present effects on human happiness; with an inquiry into our prospects respecting the future removal or mitigation of the evils which it occasions*. London: Reeves and Turner.
- Suddendorf T (2006) Foresight and evolution of the human mind. *Science* 312: 1006–1007.
- Suddendorf T, Corballis M (2007) The evolution of foresight: What is mental time travel, and is it unique to humans? *Behavioral and Brain Sciences* 30: 299–313.
- Whitehead H (2007) Learning, climate and the evolution of cultural capacity. *Journal of theoretical biology* 245: 341–350.
- Whitehead H, Richerson P (2009) The evolution of conformist social learning can cause population collapse in realistically variable environments. *Evolution and Human Behavior* 30: 261–273.
- Mesoudi A (2008) Foresight in cultural evolution. *Biology and Philosophy* 23: 243–255.
- Aiello L, Wheeler P (1995) The expensive-tissue hypothesis: the brain and the digestive system in human and primate evolution. *Current Anthropology* 36: 199–221.
- Rosenberg K, Trevathan W (1995) Bipedalism and human birth: the obstetrical dilemma revisited. *Evolutionary Anthropology: Issues, News, and Reviews* 4: 161–168.
- Bock J, Sellen D (2002) Childhood and the evolution of the human life course. *Human Nature* 13: 153–159.
- Boyd R, Richerson P (1988) *Culture and the evolutionary process*. University of Chicago Press.
- Barrett P, Gautrey P, Herbert S, Kohn D, Smith S (1987) Charles Darwin's notebooks, 1836–1844: Geology, transmutation of species, metaphysical enquiries. *British Museum (Natural History)*.
- Heath D, Heath J, Bryden C, Johnson R, Fox C (2003) Rapid evolution of egg size in captive salmon. *Science* 299: 1738–1740.
- Charnov E, Krebs J (1974) On clutch size and fitness. *Ibis* 116: 217–219.
- Lynch M (1977) Fitness and optimal body size in zooplankton population. *Ecology* 58: 763–774.
- Mangel M, Stamps J (2001) Trade-offs between growth and mortality and the maintenance of individual variation in growth. *Evolutionary Ecology Research* 3: 583–593.
- Schaffer W (1974) Optimal reproductive effort in fluctuating environments. *American Naturalist*. pp 783–790.

21. Levins R (1968) Evolution in changing environments: some theoretical explorations. Princeton University Press.
22. Slobodkin L (1964) The Strategy of Evolution. *American scientist* 52: 342–357.
23. Lewontin R (1965) Selection in and of populations. *Ideas in modern biology* 6: 299–311.
24. Ohtsuki H, Nowak M (2006) The replicator equation on graphs. *Journal of theoretical biology* 243: 86–97.
25. Petit J, Jouzel J, Raynaud D, Barkov N, Barnola J, et al. (1999) Climate and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica. *Nature* 399: 429–436.
26. Jouzel J, Waelbroeck C, Malaize B, Bender M, Petit J, et al. (1996) Climatic interpretation of the recently extended Vostok ice records. *Climate Dynamics* 12: 513–521.
27. Kimura M (1960) Optimum mutation rate and degree of dominance as determined by the principle of minimum genetic load. *Journal of Genetics* 57: 21–34.
28. Kimura M (1968) Evolutionary rate at the molecular level. *Nature* 217: 624–626.
29. Johnson T (1999) The approach to mutation-selection balance in an infinite asexual population, and the evolution of mutation rates. *Proc R Soc Lon B* 266: 2389–2397.
30. Sniegowski PD, Gerrish PJ, Johnson T, Shaver A (2000) The evolution of mutation rates: separating causes from consequences. *BioEssays* 22: 1057–1066.
31. Pease CM, Lande R, Bull JJ (1989) A model of population growth, dispersal and evolution in a changing environment. *Ecology* 70: 1657–1664.
32. Lande R, Shannon S (1996) The role of genetic variation in adaptation and population persistence in a changing environment. *Evolution* 50: 434–437.
33. Richerson P, Boyd R (2005) Not by genes alone: How culture transformed human evolution. University of Chicago Press.
34. Visser M (2008) Keeping up with a warming world; assessing the rate of adaptation to climate change. *Proceedings of the Royal Society B: Biological Sciences* 275: 649.
35. Bell G (2010) Fluctuating selection: the perpetual renewal of adaptation in variable environments. *Philosophical Transactions of the Royal Society B: Biological Sciences* 365: 87–97.