

The Evolution of Utility, Or Why Is Sex Fun?

Kevin B. Korb
Clayton School of Information Technology
Monash University
Clayton, Victoria
Australia
kbborb@gmail.com

Lachlan Brumley
Clayton School of Information Technology
Monash University
Clayton, Victoria
Australia
lachlan.brumley@gmail.com

ABSTRACT

In his 1998 book "Why Is Sex Fun?" [2] Jared Diamond assumed he was asking a rhetorical question. How could it *not* be fun for those at the tail end of billions of years of evolution? In the meantime, however, an emerging subspecies of economists have been arguing that Diamond's question needs to be taken seriously: sex is fun only by evolutionary accident, and not because it is an action promoting evolutionary fitness. They claim: 1) utilities and fitness prediction co-evolved to fill any gap in the accuracy of predictions of the reproductive value of actions ("Gap Theory"); and 2) this implies that utilities and the expected fitness of actions are distributed randomly relative to each other, and, in particular, are uncorrelated ("Non-Association Thesis"). We argue that Gap Theory has much to commend it, but that the inference to (2) is a non sequitur. Furthermore, we empirically demonstrate the falsity of Non-Association under a variety of environments realized in evolutionary simulations that satisfy Gap Theory.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Evolutionary simulation, evolution of cognition, evolution of utility, evolution of decision making, evolutionary psychology*

General Terms

Cognition

Keywords

Evolutionary simulation, evolution of cognition, evolution of utility, evolution of decision making, evolutionary psychology

1. INTRODUCTION

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '15, July 11-15, 2015, Madrid, Spain.
Copyright 2015 ACM TBA ...\$15.00.

Rational decision making, in the classic view [9], requires maximizing expected utility by, first of all, assessing as best one can the consequences of available actions and then evaluating those actions by computing their probability-weighted implications for utility. In other words, rationality is the selfish utility maximization of Equation 1

$$\max EU(a) = \sum_i U(o_i)P(o_i|a) \quad (1)$$

with a ranging over available actions and i over possible outcome states. For our purposes in this paper, we will ignore issues to do with ethics and altruism and accept this model (although, in fact, we believe ethical decision making is more rational than selfish decision making!). We shall also simply assume that cognition in general evolved, and so the ingredients of this "rational" decision making, predictive probabilities and utilities in particular, also evolved. The question we wish to consider is *how* such decision making could evolve.

Jared Diamond [2] assumed that utilities in particular must have evolved to reinforce the choice of actions that are more fit over those that are less fit, i.e., the Positive Association Thesis between utility and fitness. Those who ever found sex unfun, found extinction rational. So, whatever is left of sex is fun, while falling over cliffs is not. These sorts of claims seem not only intuitive, but intuitively *obvious*, which is exactly why his book title is funny. Rather than humor, some seem to have found a challenge in the title: to explain the evolution of utility. While we hope we retain a sense of humor, we take their challenge seriously.

This argument resides in a wider context. Just as improved understanding of evolution led to its application to psychology in the evolutionary psychology movement, evolutionary concepts have infiltrated economic thinking, leading to researchers calling themselves evolutionary economists. The group of economists concerned here take the evolution of utility seriously, however their idea of fitness is individual fitness, rather than the better justified concept of inclusive fitness of Hamilton [4] and Dawkins [1]. This is possibly due to a commitment to rational choice theory as the basis of economic analysis. While this may work in economics, individual fitness has no ability to explain the proven evolution of altruism and social cooperation, which has been at the center of research in evolutionary psychology and computational intelligence in recent decades. Individual fitness is just not fit for purpose.¹ Regardless, we can find sufficient

¹Granted, it is possible to save the *language* of individual fitness, as Nowak, Tarnita and Wilson have (in)famously

grounds for rejecting the economists’ account of the evolution of utility in their own terms, using individual fitness. For discussions of the wider concerns see, for example, Rubin and Capr [12] and Friedman [3].

2. THE GAP THEORY OF UTILITY

Two serious economists, Samuelson and Swinkels [13], begin their argument from the question “What is the value of utility?”² Their answer is consistent with recent work on the neuroscience of emotion (e.g., [11]) and also reinforcement learning theory in AI, namely that utilities serve to direct short- and intermediate-range decision making towards ultimate rewards, in their case the ultimate reward of reproduction. In reinforcement learning, actions which have immediate rewards (positive or negative) contribute some “tax” on that reward to actions in the recent past which have helped to reach the final rewarding state. In effect, intermediate actions build up a history of rewards into stable utilities associated with those actions, with the result that the more useful actions will be chosen more frequently over time. Utilities, then, are a kind of currency used to boost the choice of some actions and suppress those of others.

But, then, why would anyone need to have some actions boosted and others suppressed? If we were all perfectly rational, not merely in the sense of maximizing expected utility but also in the sense of having optimal expectations and perfect prediction, then we would evolve to always take actions that maximize our expected fitness, and there would be no role for utility to play in our decision making. The long-range effect of intermediate actions would be unambiguous, and there would be no evolutionary advantage to boosting fit actions or suppressing unfit actions. The only adaptive role possible for utilities is to fill the gap between our predictions of fitness and the systematic (evolutionarily repeated) objective probabilities of fitness outcomes in similar circumstances. Assuming utilities are strictly adaptive then, positive utilities boost the selection of actions which underestimate their objective expected fitness by the exact amount needed to compensate for the underestimation, and likewise negative utilities exactly compensate for the overestimation of fitness. We can call this the Gap Theory of Utility, and it is due to Samuelson and Swinkels. We could quibble about the fine points, but we accept it in general terms. Samuelson and Swinkels even *prove* it mathematically, for a very highly constrained set of circumstances.

While Gap Theory dictates a strict relation between utilities of actions and their expectations of fitness, it leaves entirely open the question of their correlation or lack of correlation.

2.1 Utilities and Gap Theory

argued [8]. But providing rules to translate inclusive fitness language into a far more cumbersome language of individual fitness (in effect, spreading inclusive fitness effects across all the relevant individuals) doesn’t save individual fitness as an explanatory concept. That idea is rather like suggesting that an insightful recursive algorithm written compactly in Lisp is just as insightful when translated into 1000 lines of Cobol. It just isn’t.

²Other economists have further developed the work of Samuelson and Swinkels since then, for example Rayo and Robson [10]. While the latter is interesting, it doesn’t question Samuelson and Swinkel’s main propositions and adds unnecessary complexity, so we will not treat it here.




Figure 1: A generic decision problem with utilities.

Samuelson and Swinkels present an extremely simple model of utility and decision making [13] in which agents have a choice between two actions, A and $\neg A$, which have associated probabilities of “success” – i.e., reproduction – p and q , respectively. There is a value associated with success, x , and some other value associated with failure. In order to fix a utility scale, we can simply set the value of reproduction $x = 1$. Then the remaining question is what utility should be associated with A and $\neg A$. Samuelson and Swinkels give the value of $\neg A$ by assigning 0 utiles to the combination of $\neg A$ and failure. The value of A alone is described as y and is the focus of their attention. The decision problem can then be represented by Figure 1. While their model is, of course, too simple to represent any real decision making, it nevertheless contains the needed ingredients to answer some basic questions about the evolution of utility; and the simplicity enables them their simple proof of a restricted Gap Theory.

Given this much, optimal decision making (maximizing reproductive success) is described by choosing action A if and only if $p(x+y) + (1-p)y > qx$, which is equivalent to choosing A whenever $\Delta = p - q > 0$. Since no agent lives in an ideal world, Samuelson and Swinkels introduce monotonic, continuous increasing predictive functions $\phi(s_p) = p'$ and $\phi(s_q) = q'$ [13], representing the agent’s ability to predict p and q given its perceptual state, what it can sense about the world.³ In that case, the best decision rule the agent can adopt is to choose A if and only if $p'(x+y) + (1-p')y > q'x$. Where $\delta = p' - q'$, this is equivalent to selecting A whenever $\delta + y > 0$. This latter rule is equivalent to the optimal decision whenever $\delta + y = \Delta$, which is to say, whenever the utility of A – i.e., y – exactly covers the gap between the objective probabilities p and q and the agent’s estimated probabilities $\phi(s_p)$ and $\phi(s_q)$.

Samuelson and Swinkels introduced this model of “Gap

³They claim the restrictions of monotonicity and continuity are evolutionarily plausible. This is contestable, but we shall not bother contesting it here, as the issue would be a distraction. The restriction of inputs to immediate percepts, however, is a much more serious issue than monotonicity and continuity. Samuelson and Swinkels suggest that this too is evolutionarily plausible, but in that they are mistaken. The idea that evolution would not, for example, take advantage of memory in adapting utility, which is implied, is absurd. In the future, we will develop simulation models to explore more interesting, intelligent and evolutionarily plausible predictors. In this work, however, we limit ourselves to their simpler scenario.

Theory” and indeed proved its correctness under their restrictive conditions on ϕ .

2.2 Gap Theory and Association

Samuelson and Swinkels argue, or, perhaps more accurately, *assume*, that Gap Theory implies the Non-Association Thesis. Gap Theory implies that there is evolutionary pressure for a perfect fit between the over- or under-prediction of fitness and utilities associated with that action. The Non-Association Thesis is therefore equivalent to over- and under-prediction being unrelated to the fitness of the actions being considered. So, these economists are presuming that over- and under-prediction are random accidents of evolution rather than a systematic product of evolution. In their follow-up work, Rayo and Robson present this argument, or non-argument, plainly in their abstract [10]:

Why did evolution not give us a utility function that is offspring alone? Why do we care intrinsically about other outcomes, food, for example, and what determines the intensity of such preferences? A common view is that such other outcomes enhance fitness and the intensity of our preference for a given outcome is proportional to its contribution to fitness. We argue that this view is inaccurate. Specifically, we show that in the presence of informational imperfections, the evolved preference for a given outcome is determined by the individual’s degree of ignorance regarding its significance.

It’s clear that there is a conflation of Gap Theory with the Non-Association Thesis. But, even more plainly than this, the over- and under-prediction of the fitness outcomes of our actions may well be under evolutionary control, rather than randomly distributed. Suppose, for example, that we evolved to be systematically conservative in our predictions, so that we underestimated the fitness outcomes of fitter actions and overestimated the fitness outcomes of unfit actions – i.e., that we were biased in estimating the fitness outcomes of our actions towards zero. In that case, the gap in predicting fit actions would be filled by positive utilities and the gap for unfit actions would be filled by negative utilities, yielding the Positive Association Thesis that most people believe in. The inference to Non-Association is a non sequitur so long as the premise of a random distribution of prediction error remains suppressed and unsupported. This is a surprising omission inasmuch as the large volume of work by cognitive psychologists exploring and detailing systematic human error in prediction suggests *anything but* asystematicity in the evolution of prediction error (e.g., [15, 3, 14]). Nor are Samuelson and Swinkels unaware of this literature, as they actually cite it, e.g., [16].

We can put forward what we consider a strong, positive reason, in view of Gap Theory, for believing the Positive Association Thesis and rejecting Non-Association: evolution, the engine behind evolutionary economics. In particular, we refer to the long lost prehistory of the first evolution of cognition. As we have little or no evidence about how cognition evolved from organisms that could not cognize, our considerations here are necessarily speculative, but we think our speculation is highly plausible. In particular, we suggest decision making evolved in two phases. First, organisms were driven to act by utilities, by their feelings. These

are “computed” very quickly and easily, so far as we know subjectively, and do not require memory and learning. As such, they would have evolved more readily than predictive abilities. Given their capacity for selecting fit actions, they therefore would have evolved first. Second, as some organisms began to conquer the problems of decision making in more complex environments, differential evolution would favor those who could add (or, multiply) prediction into the process.⁴ In the first phase of evolution, utilities operating without probabilities would *necessarily* be positively associated with the fitness outcomes of their actions, since, as Diamond pointed out, those lacking such an association would go extinct. Evolution, as is well known, operates on what is available to it, so the *subsequent* evolution of prediction would never have an opportunity to *disassociate* positive utility from positive fitness. Once established, Positive Association necessarily carries on, because Gap Theory demands it. So far from establishing Non-Association, given a plausible story about evolution, Gap Theory directly implies that it is false.

3. THE SIMULATED CO-EVOLUTION OF UTILITY AND PREDICTION

We shall now also demonstrate the invalidity of the argument from Gap Theory to Non-Association by describing a simulation we built that directly implements Samuelson and Swinkels’ model of decision making. In our evolutionary experiments with their model we had no difficulty finding environments which supported the evolution of all varieties of relation between utilities and the fitness of actions, positive, negative and neutral. Furthermore, in accord with the two-phase theory of its evolution, we found that when the evolution of utility precedes that of predictive ability, a positive association is invariably established and the subsequent evolution of prediction is powerless to dislodge it.

Our experimental simulation method is not the preferred tool of economics, which is rather mathematical proof; however simulation is better suited than proof to resolving many questions, including this one. Indeed, given the results of our experiments, that *any* association between fitness and utility may evolve depending upon the evolutionary circumstances, there can be no general proof of their relation.

3.1 Simulation Design

We implemented the simulation as a cellular automaton in Netlogo [17]. We used NetLogo patches as our agents, reproducing them asexually with mutation to neighboring patches (see Figure 2). We ran the simulation on a 41×41 torus world with 40 randomly generated agents to begin with. The global simulation parameters were p , p' , q , q' – that is, the probability for agents that actions A and $\neg A$ would lead to reproduction and the initial (unevolved) subjective probabilities that they used to choose which action to take – and the time steps when utilities and predictive probabilities were allowed to evolve. p and q were fixed during any one simulation run and common to all agents, while

⁴This idea, in this context, was first stated by Mascaro et al. [6]. A somewhat more complicated theory of the evolution of cognition, the Triune Brain hypothesis of MacLean [5], implies the simpler theory, but we needn’t commit to the more complicated theory.




Figure 2: The simulation in action. White patches are dead. Colored patches are agents, with the base color representing the utility associated with A (blue is more than yellow, with green in between) and the saturation the extent that the utilities cover the predictive gap (dark is poorly, light is well). This particular simulation illustrates the positive association evolved when utilities evolve first and prediction second.

p' and q' were allowed to evolve after the starting time parameter was reached in a simulation run.

Note that Samuelson and Swinkel’s function ϕ is not explicitly represented in the simulation. We may add this in future, particularly to test richer, more plausible predictive functions than they entertain. In these experiments, however, it would simply be an unneeded complication.

We found that probabilities of reproduction around 0.2 generally lead to a stable population that does not fill the world. After 300 time steps or so the great majority of simulations had reached a stable state, as reflected in evolved utilities and population size, which usually was around one half of the available patches. For the experiments we report here, we ran all simulations for 501 time steps. In order to test the two-phase evolution theory and its variations — with either utility or predictive accuracy evolving first, second, simultaneously or never — we simply set the start of their evolution to an appropriate time. (In the case of utility, there is still variability across agents even with no evolution, since each initial agent received its y from a uniform distribution centered at 0.) The result of running our simulation invariably illustrates Gap Theory, with agents having utilities that close the gap dominating and spreading through the simulated world, except in runs where the population simply dies out.

3.1.1 Agents

Each agent has the characteristics:

- A birthday and a lifespan
- An inherited y utility for action A

- p' and q' , initially set by parameter
- A record of its fitness (total offspring)

Agents choose action A if $p' - q' + y > 0$ and otherwise choose $\neg A$.⁵ The actions are successful at the rates p and q respectively. When an action leads to reproduction (i.e., is successful), the agent asexually reproduces, with Gaussian mutation applied to its utility y and predictions of success p' and q' , assuming evolution for those values is operating. The offspring is created in a randomly selected neighboring patch, unless that patch is already occupied, when the action is unsuccessful (so p and q are not the actual objective probabilities of reproduction, which in fact vary according to the population density).

In the main Netlogo display, as in Figure 2, patches are white (empty) or colored, with the base color reflecting the utility y in the range of observed utilities, with yellow representing low utility, green intermediate and blue high. Saturation reflects the extent to which the predictive gap for A has been covered by the utilities, with low saturation representing a large remaining gap and full saturation representing no gap.

3.1.2 Simulation Globals

Additional parameters are buried in the code as globals, since we had little need to adjust them once reasonable values were found. These include the Gaussian variance for mutation, the neighborhood for reproduction (one, giving us the Moore neighborhood), initial population size, and maximum lifespan.

3.2 Simulation Experiments

In all the experiments reported here, means and confidence intervals are reported for sets of 30 runs with common simulation parameters, with variation due only to the pseudo-random number generator’s seed. Since we wanted statistics from evolved equilibrium states, if a run resulted in population extinction, then another run (with a new seed taken from the clock) replaced it in the set of 30. No other selection was involved. We used Netlogo’s Behavior Space to systematically vary the simulation parameters, such as p, q, p', q' , across sets of runs.

Table 1: Mean correlations (and 95% confidence intervals) between evolved utility and fitness in three experiments, illustrating negative, neutral and positive correlations.

	(p, q, p', q')	mean correlation	CI
1	(0.15, 0.2, 0.16, 0.2)	-0.0125	± 0.0063
2	(0.2, 0.2, 0.2, 0.2)	-0.000232	± 0.00665
3	(0.2, 0.15, 0.19, 0.14)	0.01195	± 0.0074

3.2.1 Experiment 1

In our first group of experiments predictive accuracy was not allowed to evolve, while utilities were always evolving. Initially, we focused on the correlation between fitness and utility, measured over all the individuals in a single run.

⁵A good alternative is a mixed strategy, which we did. There was no interesting variation from the results of the deterministic decision making we report on here.

We were readily able to generate circumstances which induced positive, negative and neutral correlations, as shown in Table 1. Unremarkably, when predictions are accurate, utility and fitness show zero correlation; when predictions overestimate the fitness value of A , the correlation is negative; and when predictions underestimate the value of A , the correlation is positive. In the latter two cases, the mean correlations are statistically significantly distinct from zero. This is in complete agreement with Gap Theory, but violates the Non-Association Thesis. In general, we found that correlations peaked in magnitude early on in the simulation and asymptotically approached some low value thereafter, reflecting the burn-in phase when the evolutionary dynamics are strongest and an equilibrium phase when genetic diversity declines (see Figure 3).

These experiments empirically demonstrate the invalidity of inferring Non-Association from Gap Theory.




Figure 3: Evolutionary trajectory of the correlation between utility and fitness in one experimental setup (that of line 3 in Table 1).

3.2.2 Experiment 2

In this set of experiments we allowed the evolution of both utility and subjective probability. We looked at evolution when the initial predictions were over- and under-confident about the fitness of A and when they were accurate, and allowing utility to evolve first, second, or simultaneously with predictive ability. This allowed us to test directly our two-phase evolution theory — that an initial evolution of utility followed by an evolution of predictive accuracy would result in a positive correlation between fitness and utility. The fitness prediction for $\neg A$ was always initially accurate ($q = q' = 0.1$). The results in evolved utility for A (i.e., y) are shown in Table 2. The cases corresponding to our two-phase evolution theory are in the top third (where predictions are initially underconfident), and these show statistically significantly higher evolved utilities for comparable runs where utility evolves earlier (e.g., comparing line 7 with line 1) and for comparable runs where the predictions are accurate or too high (e.g., comparing line 7 with lines 16 and 25). The same data are shown as a histogram in Figure 4, but clustered according to the time when evolution of prediction and utility was allowed to begin. The evolved utility is clearly greater when utility evolves first as well as when the initial prediction underestimates the fitness of A (the red bars).

Table 2: Mean evolved utilities with 95% confidence intervals for Experiment 2. Results are listed in groups of 9 sets of 30 runs each. The top set has initial under-prediction of the fitness of A ; the next starts with accurate prediction; the third set initially over-predicts fitness.

	prob-start	utility-start	mean utility	CI
1	0	0	0.630	± 0.069
2	0	150	0.402	± 0.038
3	0	520	0.229	± 0.021
4	150	0	0.838	± 0.043
5	150	150	0.561	± 0.033
6	150	520	0.266	± 0.018
7	520	0	0.876	± 0.036
8	520	150	0.705	± 0.023
9	520	520	0.231	± 0.016
10	0	0	0.583	± 0.051
11	0	150	0.379	± 0.050
12	0	520	0.188	± 0.028
13	150	0	0.735	± 0.040
14	150	150	0.519	± 0.030
15	150	520	0.225	± 0.022
16	520	0	0.737	± 0.031
17	520	150	0.622	± 0.026
18	520	520	0.186	± 0.021
19	0	0	0.524	± 0.052
20	0	150	0.363	± 0.033
21	0	520	0.159	± 0.022
22	150	0	0.660	± 0.034
23	150	150	0.462	± 0.037
24	150	520	0.177	± 0.026
25	520	0	0.721	± 0.033
26	520	150	0.527	± 0.027
27	520	520	0.138	± 0.020

4. CONCLUSION

We have argued and, with our simulation, successfully demonstrated that Gap Theory does not imply any particular association between evolved utilities and fitness. Not implying a particular association is not the same as implying no association, of course. Those claiming otherwise, partly for the shock value of asserting the Non-Association Thesis, have made elementary mistakes, apparently unintentionally conflating Gap Theory with Non-Association and so not noticing the inferential gap in their position. We have, on the other hand, given some modest reason to think that there should be a positive association between utility and fitness and, somewhat unusually for cognitive psychology, confirmed common intuition. Our simulation supports our claims.

This kind of work in the simulated evolution of cognition has not been done before, but potentially has a large role to play in advancing our understanding of cognition. Ongoing work in evolutionary psychology, and its application to economics and politics (as in the work of McDermott et al. [7] and Rubin and Capra [12]), shows that the utilities revealed in many human behaviors have deep origins in the early evolution of homo sapiens. We would argue that these origins are even deeper, lying in the early evolution of *animals* in general. Richer simulations, in particular simulations which allow for a more detailed representation of the co-evolution

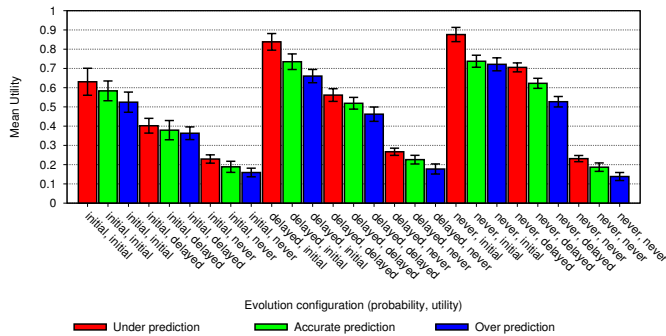


Figure 4: Histogram of the data in Table 2, with experiments reordered according to the evolution start times of prediction and utility: the first third are those with prediction evolving first (initial evolution, with varying times of onset for utility evolution); the second third with prediction evolution delayed; the third with prediction not evolving.

of utility and decision making, can support, or refute, claims such as ours. We could, and, we hope, will, learn a lot about the evolution of cognition by studying it *in silico*.

Acknowledgements

We acknowledge the assistance of Steven Mascaro in discussions about this paper, as well as the ANU School of Philosophy and the Evolutionary Computation Research Group at Victoria University, Wellington.

5. REFERENCES

[1] R. Dawkins. *The Selfish Gene*. Oxford University Press, 1976.

[2] J. Diamond. *Why Is Sex Fun?* Basic Books, 1998.

[3] D. Friedman. Economics and evolutionary psychology. In R. Koppl, editor, *Evolutionary Psychology and Economic Theory*, volume 7 of *Advances in Austrian Economics*, chapter 2, pages 17–34. Elsevier, 2004.

[4] W. D. Hamilton. The genetical evolution of social behavior, I and II. *Journal of Theoretical Biology*, 7:1–16 & 17–52, 1964.

[5] P. MacLean. *The Triune Brain in Evolution: Role in Paleocerebral Functions*. Springer, 1990.

[6] S. Mascaro, K. B. Korb, A. E. Nicholson, and O. Woodberry. *Evolving Ethics: The New Science of Good and Evil*. Imprint Academic, 2011.

[7] R. McDermott, J. H. Fowler, and O. Smirnov. On the evolutionary origin of prospect theory preferences. *The Journal of Politics*, 70:335–350, 2008.

[8] M. A. Nowak, C. E. Tarnita, and E. O. Wilson. The evolution of eusociality. *Nature*, 466:1057–1062, 2010.

[9] F. P. Ramsey. Truth and probability. In R. Braithwaite, editor, *The Foundations of Mathematics and Other Essays*. Humanities Press, New York, 1931.

[10] L. Rayo and A. Robson. Biology and the arguments of utility. Cowles Foundation Discussion Paper, 2014.

[11] E. Rolls. *Emotion Explained*. Oxford University Press, 2005.

[12] P. H. Rubin and C. M. Capr. The evolutionary psychology of economics. In S. C. Roberts, editor, *Applied Evolutionary Psychology*, pages 7–15. Oxford University Press, 2011.

[13] L. Samuelson and J. Swinkels. Information, evolution and utility. *Theoretical Economics*, 1:119–142, 2006.

[14] K. E. Stanovich and R. F. West. Evolutionary versus instrumental goals: How evolutionary psychology misconceives human rationality. In D. Over, editor, *Evolution and the Psychology of Thinking: The Debate*, pages 171–230. Psychology Press, Hove, England, 2003.

[15] A. Tversky and D. Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185:1124–1131, 1974.

[16] A. Tversky and I. Simonson. Context-dependent preferences. *Management Science*, 39:1179–1189, 1993.

[17] U. Wilensky. Netlogo, 1999. <http://ccl.northwestern.edu/netlogo/>.