THE EVOLUTIONARY FOUNDATIONS OF PREFERENCES*

Arthur J. Robson
Department of Economics
Simon Fraser University
8888 University Drive
Burnaby, BC, Canada V5A 1S6
Robson@sfu.ca

Larry Samuelson
Department of Economics
Yale University
30 Hillhouse AVenue
New Haven, CT 06520-8281, USA
Larry.Samuelson@yale.edu

April 12, 2010

Abstract. This paper, prepared for the forthcoming *The Social Economics Handbook* (Jess Benhabib, Alberto Bisin and Matthew Jackson, editors, Elsevier Press), surveys recent work on the evolutionary origins of preferences. We are especially interested in the circumstances under which evolution would push preferences away from the self-interested perfectly-rational expected utility maximization of classical economic theory in order to incorporate environmental or social considerations.

---

*We thank Andrew Postlewaite for helpful discussions and comments. Arthur Robson thanks the Canada Research Chair Program and the Social Sciences and Humanities Research Council of Canada and Larry Samuelson thanks the National Science Foundation (SES-0549946 and SES-0850263) for financial support.

The Evolutionary Foundations of Preferences

by *Arthur J. Robson and Larry Samuelson*

Contents

The Evolutionary Foundations of Preferences

by *Arthur J. Robson and Larry Samuelson*

1 Introduction

This essay on the evolutionary foundations of preferences is best introduced with an example. The example in turn requires some notation, but this seemingly technical beginning will set the stage for an ensuing discussion that is more intuitive.

We are interested in a setting in which consumption must be distributed across periods in the face of uncertainty. Suppose that time is discrete, indexed by $\{0, 1, 2, \ldots\}$. A state $\omega \in \Omega$ is first drawn from the finite set Ω, with $\rho(\omega)$ giving the probability of state ω. The consumption bundle in period t is drawn from a set C and given by $c_t(\omega)$, being a function of the period and the realized state. The consumption profile, identifying a consumption bundle for each period, is then $\{c_t(\omega)\}_{\omega\in\Omega, t\in\{0,1,\ldots\}}$. Let c denote a typical such consumption profile and $\mathcal{C}$ the set of such profiles. How do we model preferences over the set $\mathcal{C}$?

The most common approach in economics is to assume there exists an increasing utility function $u : C \rightarrow \Re$, allowing preferences over $\mathcal{C}$ to be represented by the discounted-sum-of-expected-utility function $U : \mathcal{C} \rightarrow \Re$, given by

$$U(c) = \sum_{t=0}^{\infty} \sum_{\omega\in\Omega} D^t u(c_t(\omega))\rho(\omega), \tag{1}$$

where $D \in (0, 1)$ is the discount factor. Dating at least to Samuelson [135], this model is so familiar as to require no explanation and no second thoughts when pressed into service.

Why is this a useful representation? From an analytic point of view, (1) is compelling for its tractability. The additive separability across time and states, the stationarity of the discounting, and the stationarity of the function u over time and states all make analysis and computation easier. For example, this maximization problem exhibits the consistency property that lies at the heart of dynamic programming. Computationally, a single function u is much easier to simulate or estimate than one such function for each period or state. At the very least, one might view (1) as an ideal point of departure for a study of behavior, however unrealistic it turns out to be, perhaps with the goal of subsequently examining the robustness of its more interesting implications to more flexible specifications.

From a normative point of view, (1) can be viewed as an expression of rationality. Within periods, the expected utility formulation is implied by Savage's [136] axioms, often defended as foundations of rationality (with Allais [4] and Ellsberg [38] giving rise to a vast literature questioning their positive applicability). For example, a person whose behavior is characterized by (1) can never fall prey to a money pump, a criterion typically regarded as essential for rationality (cf. Nau and McCardle [96]). Looking across periods, it is once again reassuring that the resulting behavior is consistent, in the sense that an optimal consumption plan at time t is the continuation of the optimal plan at time $t' < t$. This ensures that recommendations based on (1) cannot lead to conflicting advice.

From a positive point of view, however, (1) is less convincing, doing both too little and too much. This representation does too little in the sense that it leaves important questions open. What is the shape of the function u? Are people risk-neutral, risk-averse, risk-seeking, or something more complicated? How are risk attitudes related to observable characteristics of either the decision maker or her environment? The representation does too much in the sense that it places a great deal of structure on preferences. Do people really discount in such a stationary fashion? Are their preferences linear in probabilities? Do they think in terms of probabilities at all? Are their preferences really so separable? Once we go beyond these points to open the deeper question of what enters the utility function, all sorts of questions arise. Are people really concerned only with their own consumption and nothing else? How might their preferences be affected by various aspects of their environment, including perhaps the consumption of others?

One possible response to these questions is empirical. Bolstered by ever-more-plentiful data as well as more powerful experimental techniques, we can simply observe behavior and infer the corresponding preferences. In doing so, one could usefully draw on the rich revealed-preference literature in psychology as well as economics.[1]

Our thinking on this point is that empirical work on preferences and behavior is essential. However, the specification of preferences is sufficiently complicated, and poses sufficient identification issues, that we have little hope of making progress by pursuing a *purely* empirical approach. However much data we have, we can hope to make sense of it only in the context of theoretical models.[2] But where do we find these models? Building models is

---

[1] See Rabin [110] for an introduction to work at the intersection of psychology and economics.

[2] See Gilboa and Samuelson [60] for an abstract discussion of this point.

something at which economists excel, and economists are seldom idle when there are new models to be produced. As one might expect, the analysis of of preferences is no exception.[3] The difficulty is that if we do not restrict ourselves to some simple form such as (1), it seems that anything goes, and we can provide theoretical foundations for anything. How do we impose discipline on the resulting theoretical exercise?

This quest for discipline is perhaps the ultimate motivation for (1). Whatever its disadvantages, it clearly imposes a great deal of structure on the analysis. As a result, when faced with behavior seemingly inconsistent with (1), a common reaction is to preserve (1) while searching for features of the environment to account for the proposed behavior. Postlewaite [109] states the case for doing so quite clearly. By allowing departures from (1) as explanations, not only may we acquire sufficient explanatory power as to rob the resulting exercise of any substance, but the ease with which we can thereby accommodate observed behavior may distract attention from aspects of the environment that actually lie behind the behavior. If allowed to work freely with models in which people simply prefer to not purchase used durable goods such as automobiles, we may never have discovered the lemons phenomenon (Akerlof [2]). It may thus be better to stick with (1), trading the constraints it imposes and its potential lack of realism for the concreteness it brings to our inquiry.

The point of departure for this essay is the belief that we must both sometimes impose more structure on (1) as well as sometimes move beyond this formulation, and that we require solid theoretical foundations for both. We suggest seeking the required theoretical discipline in evolutionary models. In particular, we view human preferences as having been shaped by years of evolutionary selection. When thinking about whether (1) is a reasonable representation of preferences, or which more specific or more general models might be useful alternatives, our first step is to ask what sorts of preferences are likely to emerge from this evolutionary process. The more readily can we provide evolutionary foundations for a model of preferences, the more promise we see in using this model in theoretical and applied economic analyses.

This approach to preferences raises a collection of methodological issues that are discussed in Section 2. Sections 3 and 4 provide illustrations from the literature. Section 3 concentrates on the functional form assumptions built into (1), including the expected-utility criterion that is applied within periods and the exponentially-discounted summation that aggregates utility

---

[3] Camerer, Loewenstein and Rabin [22] provide a good point of entry into this literature.

across periods. Section 4 examines arguments that are likely to appear in the utility function beyond an agent's own consumption. Section 5 very briefly concludes.

2 Evolutionary Foundations

2.1 Evolution and Economic Behavior

Is it reasonable to talk about evolution and human behavior at all? A large literature, referred to as evolutionary game theory, has grown around evolutionary models of behavior.[4] The presumption behind evolutionary game theory is that human behavior, whether in games (and hence the name) or decision problems, typically does not spring into perfect form as the result of a process of rational reasoning. Instead, it emerges from a process of trial and error, as people experiment with alternatives, assess the consequences, and try new alternatives. The resulting adaptive processes have been modeled in a variety of ways, from Bayesian to reinforcement learning, from cognitive to mechanical processes, from backward to forward looking processes, all collected under the metaphor of "evolutionary game theory."

This literature has provided valuable insights into how we interpret equilibria in games, but we have a fundamentally different enterprise in mind when talking about the evolution of preferences in this essay. We take the word "evolution" literally to mean the biological process of evolution, operating over millions of years, that has brought us to our present form.[5] The driving force behind this evolution is differential survival and reproduction. Some behavior makes its practitioners more likely to survive and reproduce than others, and those behaviors most conducive to survival are the ones we expect to prevail. Our task is to identify these behaviors.

This view would be uncontroversial if we were talking about the evolution of physical characteristics. A giraffe who can reach more leaves on a tree is

---

[4] See, for example, Fudenberg and Levine [55], Hofbauer and Sigmund [73], Mailath [89], Samuelson [131], van Damme [155, Chapter 9], Vega Redondo [157], Weibull [159] and Young [167].

[5] We have no doubt that cultural evolution is also vitally important. We expect the techniques we examine to transfer readily to models of cultural evolution, often with simply a reinterpretation. We find interpretations in terms of biological evolution more straightforward, and hence tend to adopt them. Henrich, Boyd, Bowles, Camerer, Fehr, Gintis and McElreath [72] and Henrich, Boyd, Bowles, Camerer, Fehr and Gintis [71] provide interesting points of departure into the study of cultural evolution and economic behavior.

more likely to survive, and hence evolution gives us giraffes with long necks. A bat that can detect prey is more likely to survive, and so evolution gives us bats capable of echolocation. Porcupines are more likely to survive if they are not eaten, and so have evolved to be covered with sharp quills. The list of such examples is virtually endless.

Behavior can also confer an evolutionary advantage, with a similarly long list of examples. African wild dogs enlarge their set of eligible prey, and hence their chances of survival, by hunting in packs. Vampire bats reduce their likelihood of starvation by sharing food. Humans enhance the survival prospects of their offspring by providing food for their young. If different members of a population behave differently, then those whose behavior enhances their survival can be expected to dominate the population. The relentless process of differential survival will thus shape behavior as well as characteristics.

Doesn't this commit us to a strong form of biological determinism? Is our behavior really locked into our genes? We think the answer is no on both counts.[6] Nature alone does not dictate behavior. However, there is a huge gap between the assertion that genetic factors determine every decision we will ever make and the assertion that biological considerations have no effect on our behavior. We need only believe that there is some biological basis for behavior, however imprecise and whatever the mechanics, for the issues raised in this essay to be relevant.[7]

2.2 The Rules of Evolution

We will often refer to "evolution" as if referring to a conscious being. We will use phrases such as "evolution selects" or "evolution prefers" or "evolution maximizes" or even "evolution believes." It is important to be clear at the beginning that we attribute no consciousness and no purpose to evolution. We have in mind throughout the standard, mindless process of mutation and selection studied by biologists. We suppose that individuals in a population may have different types, whether these are manifested as different physical characteristics or different behavior. These different types reflect genetic endowments that arose initially from undirected, random mutations.

---

[6] Ridley [116] provides an accessible introduction to the voluminous literature that has grown around these sometimes controversial questions.

[7] The evidence that there is some such connection is both wide-ranging and fascinating. For two examples, see Dreber and Hoffman [33] and Knafo, Israel, Darvasi, Bachner-Melman, Uzefovsky, Cohen, Feldman, Lerer, Laiba, Raz, Nemanov, Gritsenko, Dina, Agam, Dean, Bronstein, and Ebstein [78].

Some of these types will make their possessors more likely to survive, others will be detrimental. Over time, this process of differential survival will cause a larger proportion of the population to be characterized by the former types, and it is this process that lies behind our analysis.[8] If allowed to run unchecked, the pressures of differential survival will eliminate those types that are less likely to survive and produce a population consisting only of those whose behavior is most conducive to survival. As a result, it is often convenient to model the *outcome* of an evolutionary process as the solution to a maximization problem. This convention is familiar to economists, who routinely model consumers, firms, governments, and other entities as maximizers, bolstered by the view that this maximization may be the outcome of an adaptive process rather than conscious calculation. We proceed similarly here when talking about evolution, without any illusions that there is purposeful behavior behind this maximization.

The idea that an evolutionary perspective might be helpful in studying behavior is by no means unique to economists. The field of evolutionary psychology has grown around this view of behavior.[9] We can learn not only from the successes of evolutionary psychology, but also its difficulties. Gould and Lewontin [61] criticize evolutionary psychology as being an exercise without content. In their view, a clever modeler can produce an evolutionary model capable of producing any behavior. To reinforce their point, they refer to the resulting models as "just-so" stories. As we have already noted, of course, an analytical approach capable of explaining everything in fact explains nothing. If an evolutionary approach is to be useful, we must address the just-so critique.

Economists are also adept at constructing models, and the criticism that we can concoct models rationalizing any imaginable sort of behavior is not a new one. How do we reconcile Gould and Lewontin's argument with our assertion that evolutionary models are designed to impose discipline on our study of preferences? In our view, the ability to fix a characteristic of behavior and then construct an evolutionary rationale for that behavior is only the first step. If we can go no further, we have typically learned very little. An obvious next step is to fit the model into its place in the existing body of evolutionary theory. Simple and direct models constructed from familiar and inherently-plausible evolutionary principles tend to be convinc-

---

[8] We suggest Dawkins [32], Ridley [115] and Williams [160] as accessible introductions to evolutionary theory, and Hofbauer and Sigmund [73] for a more precise examination of the conditions under which the outcome of an evolutionary process can be modeled as the solution to an optimization problem.

[9] Barkow, Cosmides and Tooby [9] provide a wide-ranging introduction.

ing, while convoluted models taking us well beyond the usual evolutionary considerations are reasonably greeted with skepticism. Moving beyond this informative but subjective evaluation, our goal should be to construct models that generate predictions beyond those of the target behavior, especially predictions that we could take to data. The more fruitful is a model in doing so, the more useful will it be.

2.3 Evolution and Utility Functions

The preceding subsections have referred frequently to the evolution of behavior, while our title refers to the evolution of preferences. How should we think about evolution shaping our behavior? In one view, evolution would simply program or "hard-wire" us with behavior, equipping us with a rule indicating what to do in each possible circumstance. Alternatively, we might think of evolution as equipping us with utility functions and instructions to maximize utility whenever called upon to make a choice. Most of what we discuss in this essay requires no choice between these alternatives, and requires us to take no stand on the countless intermediate constructions that combine aspects of both types of model. Our focus will primarily be to identify *behavior* that confers evolutionary advantage. We will then frequently describe this behavior in terms of the preferences with which it is consistent. However, this description is a matter of convenience rather than an assertion about causality.

Taking this approach keeps us squarely within the revealed-preference approach to behavior. Among the fundamental building blocks of economic theory is an assumption that behavior satisfies the consistency conditions captured by the revealed-preference axioms. However, it is often insightful to describe this behavior in terms of preferences, and then convenient to use these preferences as the point of departure for subsequent models of behavior. Similarly, it is behavior that matters to evolution, but there often will be much to be gained by describing this behavior in terms of preferences.[10]

No amount of introspection will tell us the extent to which our behavior is hard-wired and the extent to which we have discretion. Reading a restaurant menu and choosing a meal makes us feel as if we have conscious control over our actions. However, there is no particular reason why that same feeling could not accompany an inevitable action, or why we might not make choices

---

[10] This emphasis on behavior as the primitive object of analysis distinguishes the evolutionary approach from much of behavioral economics, where the process by which choices are made often takes center stage. See Camerer [21] and Gul and Pesendorfer [65] for a discussion of these issues.

without being aware of what we are doing. Pursuing these distinctions runs the risk of recreating a long-running discussion of whether we have free will, and how we would know whether we have. This is a fascinating topic, but one that has bedevilled philosophers for centuries and that would only be a hopeless diversion here.

At the same time, we think there are good a priori grounds for thinking of evolution as designing us to be utility maximizers rather than simply hard-wiring us with behavior, and Section 4.2.2 relies on a view of utility maximization as a process that shapes our choices. Robson [122] offers an argument for the evolutionary utility of utility functions, beginning with the assumption that environments fluctuate more quickly than evolution can respond. Simply telling people to hunt rabbits is risky because they may encounter situations in which deer are more readily available. With hard-wired behavior, an evolutionary response to such situations would require a deer-hunting mutation, or perhaps several if the first few such mutations are unlucky. This must then be followed by a process of selection that may be fast compared to length of time humans have been developing, but may be quite slow compared to the length of time it takes for a shock to the weather or to the population of predators to once again make rabbits relatively plentiful. By the time the new hard-wired behavior has spread into the population, it may well be out of step with the environment. A more flexible design would give the agent the ability to observe and collect information about her environment, coupled perhaps with an instruction of the form "hunt the relatively more plentiful prey." This type of contingent behavior will be effective as long as evolution can reasonably anticipate the various circumstances the agent may face. However, this may require taking account of a list of contingencies prohibitively long for evolution to hit upon via trial-and-error mutations. A more effective approach may then be to endow the agent with a goal, such as maximizing caloric intake or simply feeling full, along with the ability to learn which behavior is most likely to achieve this goal in a given environment. Under this approach, evolution would equip us with a utility function that would provide the goal for our behavior, along with a learning process, perhaps ranging from trial-and-error to information collection and Bayesian updating, that would help us pursue that goal.[11]

---

[11] There are, of course, other aspects of our preferences that evolution may prefer to place outside our learning. Many people have a deep-seated fear of snakes (cf. Mineka and Cook [94] and Pinker [106, pp. 388–389]), but few of us are afraid of mushrooms. Since both can be potentially fatal and both can be eaten, this combination is by no means obvious. To see why we may have come to such a state, imagine that being bitten by a

If this is the case, however, why would we attach utility to activities such as eating? Evolution necessarily selects for that behavior which leads to the most effective propagation, so why don't we have preferences solely over offspring, or some appropriate trade-off between the quantity and quality of offspring, or some other measure of descendants? One difficulty is that simply giving us preferences over offspring gives rise to a small-sample learning process. Human offspring come relatively rarely and provide relatively sparse feedback. Opportunities to eat are much more frequent and provide a much richer flow of information. An agent designed with the goal of producing healthy adult offspring, and then left to learn the details of doing so by trial-and-error, may not learn soon enough to do any good. An agent whose goal is to be well nourished may acquire enough experience soon enough to make good use of this information. Defining utilities in terms of offspring thus gives us an objective that quite faithfully captures the relevant evolutionary criterion, but gives us little means of learning how to accomplish this objective. Defining utilities in terms of intermediate goods such as consumption gives us an objective that only approximates evolution's—in some environments we will mechanically pursue additional consumption even though circumstances are such that doing so retards reproduction—in return for giving us the means to effectively learn how to accomplish this objective. The choice of which arguments to place in a utility function thus reflects a delicate evolutionary balancing act, one that we believe merits further study. As a first step, there is much to be learned about this evolutionary tradeoff simply from observing how evolution has solved this problem, i.e., observing what enters our utility functions.

Utility functions carry risk for evolution as well as benefits. Evolution has equipped us with preferences over many things—basic needs, such as food, sleep, safety, sex, and shelter, as well as more complicated items such as our relationship with others and our position in our community—that evolution has chosen because of the resulting salutary effects on our fitness. The fact that we have cognitive abilities that allow us to predict the effects of our actions, and to choose actions whose effects fare well in terms of our preferences, suggests that the resulting behavioral flexibility is also

---

poisonous snake is very unlikely to happen but likely to be fatal if it does, while ingesting a poisonous mushroom is more likely to occur but less likely to be fatal. Then evolution may optimally leave it to her agents to sort out which mushrooms are dangerous, while being unwilling to take chances on encounters with snakes. In general, evolution should make us fear not simply things that are bad for us, but rather things whose danger we may underestimate *without* discovering our error before they kill us. Samuelson and Swinkels [134] pursue these possibilities.

evolutionarily advantageous. At this point, however, a conflict can arise between evolution's preferences and our preferences. We have been designed to maximize our utility or "happiness," while evolution does not care whether we are happy, instead viewing happiness simply as a means for producing evolutionarily valuable ends. Maximizing happiness must on average lead to good evolutionary outcomes, or our utility functions would be designed differently, but this still leaves room for conflict. Evolution has given us a taste for sex, but over the course of having children we may notice some of the sometimes less desirable effects, leading to birth control practices that can thwart evolution's goals. It is important to bear the potential for such conflict in mind when confronted with behavior that seems otherwise inexplicable.

2.4 Evolutionary Mismatches

There are two complementary approaches to thinking about the evolutionary foundations of behavior. One is based on the observation that we currently live in an environment much different from that in which we evolved. As a result, behavior that was well suited for our evolutionary environment may fit quite awkwardly into our current one. For example, food was likely to have been in perpetually tenuous supply over the course of our evolutionary history, and the only technology for storing it was to eat it. An instruction of the form "eat all you can whenever you can" accordingly may have made good evolutionary sense. This presumably explains why so many of us struggle to keep our weight down in our modern world of abundance. Similarly, predators were probably not only a threat during much of our evolutionary history, but one that often left little leeway for learning. Ascertaining which animals are dangerous by trial-and-error is a process fraught with danger, even if most animals pose no threat. A deep-seated fear of predators was accordingly quite useful for survival. This presumably explains why children in our modern urban society are much more likely to fear wild animals than electrical outlets, even though the latter pose a much greater threat.

We refer to these types of observations as "evolutionary mismatch" models. This is clearly a useful perspective.[12] However, our interest will typically lie not in such mismatch stories, but in examining behavior that is well adapted to its environment. We will accordingly be especially interested in tracing various features of behavior to features of the environment in which the behavior could have evolved. For example, we will examine

---

[12] See Burnham and Phelan [20] for a wealth of examples.

how the nature of the uncertainty in the environment affects intertemporal preferences. Mismatches are clearly important, but we believe that a good understanding of how preferences are tailored to the environment in which they evolved is an essential first step in understanding their effects in mismatched environments. If nothing else, allowing ourselves to indulge in mismatch explanations gives us yet one more degree of freedom in constructing our models, while the goal throughout is to use evolutionary arguments to restrict such freedom.

It is important throughout to distinguish evolutionary mismatches from the potential conflict, noted in Section 2.3, between evolutionary goals and the results of our utility maximization. The latter conflict readily arises in the environment in which we evolved. Evolution finds it expedient to give us utility functions because it is prohibitively difficult to simply dictate every aspect of our behavior. But once this step is taken, the prospect arises that the resulting utility maximization will sometimes lead to counterproductive outcomes, even before we consider the effects of thrusting the agent into a new environment.

2.5 The Indirect Evolutionary Approach

We distinguish the work described in this essay from a body of literature that has come to be called the "indirect evolutionary approach." It is worth making this distinction carefully. The indirect evolutionary approach grew out of evolutionary game theory. In the simplest evolutionary-game-theory model, players are characterized by the actions they take in the decision problem or game of interest. We might think of the players as being programmed to take such actions. As play progresses, a revision protocol induces a process by which the players switch their actions. For example, players may randomly choose a new action whenever their realized payoff falls below an aspiration level, or players may switch after each period to the action that would have been a best response to the previous-period average population action, or may switch only in randomly-drawn periods to actions that are best responses to an average of the play of their previous opponents, and so on. One can imagine an endless list of such revision protocols. A central question in evolutionary game theory concerns the extent to which we can characterize the outcome of such revision protocols over the course of repeated play. Will the people be directed to behavior that appears to be "rational?" For example, will their behavior satisfy the revealed preference axioms? Will it maximize a simple objective? Will people eschew dominated strategies? Will the process induce population behavior that can be rationalized by a concept such as Nash equilibrium? Will the resulting behavior satisfy more refined equilibrium concepts?

The point of departure for the indirect evolutionary approach is to note that throughout the rest of economics, we typically model people as being characterized by preferences rather than simply actions, with these preferences inducing actions through a choice procedure such as utility maximization. Taking this idea literally in an evolutionary context, we can think of people as maximizing utility given their preferences, with their preferences adjusting over time according to a revision protocol. The evolutionary process now shapes behavior through its effect on preferences, and it is this indirect link that gives rise to the name indirect evolutionary approach, pioneered by Güth [66] and Güth and Yaari [67].

The indirect evolutionary approach has been embraced by many because of its ability to explain seemingly anomalous preferences. To see what is involved, it is useful to start with an example. Consider the following game:[13]

$$\begin{array}{c|c|c|}
 & L & R \\ \hline
T & 6,2 & 4,4 \\ \hline
B & 5,1 & 2,0 \\ \hline
\end{array} \quad . \tag{2}$$

This game has a unique Nash equilibrium, given by (T, R), with payoffs $(4, 4)$.[14]

Now suppose that, before the game begins, player 1 could commit to playing B, and that player 2 can observe whether such a commitment has been made. The game proceeds as before if no commitment is made, and otherwise player 1 is locked into B and 2 is left to choose an action. Essentially, a commitment gives us a new game with a sequential structure in which player 1 moves first. This new structure is valuable for player 1. By committing to B, 1 can ensure player 2 will choose a best response of L, giving player 1 a payoff of 5. It is clear that player 1 would jump at the chance to commit.

The observation that commitments can be valuable has a long history, beginning with Stackelberg ([158], translated into English in Peacock [104]) and playing a prominent role in Schelling [137]. Early theories of bargaining, including Binmore [15] and Crawford and Varian [29], explore the power of commitment more formally, as does Frank [45]. While it is straightforward to see that it can be valuable to make commitments, it is less clear just how one does so.

---

[13] The subsequent discussion follows Samuelson [132],

[14] This is the unique rationalizable outcome, since the strategy T strictly dominates B and R is a strict best response to T.

Now let us think of a population of player 1s and another population of player 2s. Players from these populations are repeatedly matched to play the game given by (2). The indirect evolutionary approach assumes that the payoffs in (2) are "material payoffs" or "fitnesses." These are the payoffs that are relevant in evolutionary terms. Evolution induces behavior by endowing agents with preferences over the actions T and B (for player 1s) and L and R (for player 2s). These preferences need not match the fitnesses given in (2), but it is fitnesses and not preferences that govern the evolutionary process. Agents whose behavior leads to high fitnesses will reproduce relatively rapidly and the population will ultimately be dominated by such preferences. In particular, an agent may choose an action that performs wonderfully from the point of view of the agent's preferences, all the while wasting away in the population because the action yields a low fitness. Evolution can thus mislead her agents, in the sense that preferences need not match fitnesses, but cannot fool herself, in that high fitnesses remain the ticket to evolutionary success.

Is there any reason for preferences to be anything other than fitnesses in such a setting? The key here is the assumption that preferences are observable, in the sense that when two players meet, each player can observe the other's preferences. The two matched players then play a complete-information version of the game given by (2), with their behavior governed by their preferences, and with the evolutionary implications of their behavior governed by the fitnesses given in (2). Suppose that player 2s have preferences that match fitnesses, as do some player 1s. However, the population also includes some player 1s whose preferences make B a strictly dominant strategy, effectively committing themselves to B. In response to the former types of player 1, player 2 will choose R, giving 1 a payoff of 4. In response to the latter, player 2 will choose L, giving 1 a payoff of 5. As a result, the population will eventually be dominated by player 1s committed to playing B. There is thus evolutionary value in equipping agents with preferences that do not reflect their fitnesses.

Bolstered by results such as this, the indirect evolutionary approach has been interpreted as providing foundations for a collection of empirical, experimental, or introspective findings that appear inconsistent with material self interest, including the endowment effect, altruism, vengeance, punishment, and so on.[15] These results are intriguing, but raise two questions. First, initial applications of the indirect evolutionary approach typically considered only a few possible preference specifications, often including preferences that

---

[15] See Ostrum [103] for an introduction.

match material fitnesses and one or more "commitment preference" alternatives that are tailored to the game in question. In considering (2), for example, we considered the possibility that 1 might be committed to B, but there are many other possible preference specifications. What happens if they are present as well? Player 2, for example, would like to commit to R, for much the same reason that 1 finds it valuable to commit to B. What if there are also player 2s who are so committed? What if the entire collection of preference specifications were allowed? Would we be confident that the commitment types emerging from simple models would also be selected from such a crowd?

More importantly, it was critical in the preceding argument that players could observe each other's preferences. Being committed to B is an advantage to player 1 only because it affects player 2's behavior, inducing 2 to switch to L. Ely and Yilankaya [40] and Ok and Vega Redondo [102] confirm that if preferences are not observable, any limit of behavior in their indirect evolutionary models must constitute a Nash equilibrium in material fitnesses. The indirect evolutionary approach with unobservable preferences then gives us an alternative description of the evolutionary process, one that is perhaps less reminiscent of biological determinism, but leads to no new results.

Preferences are not typically high on the list of things taken to be observable in economic analysis. Is it reasonable to assume that people can identify one another's preferences? Frank [46] argues that we do often have good information about the preferences of others, and that there is a technological basis for such information. Our preferences are determined partly by emotions such as anger or embarrassment that are beyond our conscious control, expressed by involuntary changes in our facial expressions and body language. If one is prone to blushing when the center of attention, how much good does it do to remind oneself not to blush? Who can keep flashes of anger out of their eyes? Our preferences may then often be an open book free for others to read. At the same time, Güth [66] shows that preferences need not be *perfectly* observable in order for the indirect evolutionary approach to have nontrivial implications. It suffices that player 2 *sometimes* be able to discern player 1's preferences and react to them. As Güth notes, it is a seemingly quite strong assertion that this is never the case, arguably as unrealistic as the assumption that people can always observe one another's preferences.

To evaluate these considerations, we must return to the evolutionary context. The standard argument is that we can observe preferences because people give signals—a tightening of the lips or flash of the eyes—that provide clues as to their feelings. However, the emission of such signals and their correlation with the attendant emotions are themselves the product of evolution. A complete version of the indirect evolutionary approach would then incorporate within the model the evolution of preferences *and* the evolution of the attendant signals. In (2) for example, player 1 prefers (T, L) to (B, L). Evolution thus has an incentive not only to produce player 1s who are visibly committed to playing B, but also a version of player 1 whose signals match those emitted by those player 1s committed to B, inducing L from player 2, but who then plays T. What prevents the appearance of such a mimic? We cannot simply assume that mimicry is impossible, as we have ample evidence of mimicry from the animal world, as well as experience with humans who make their way by misleading others as to their feelings, intentions and preferences.[16] If such mimics did appear, of course, then presumably player 2s would at least eventually learn that player 1s appearing to be committed to B are not always so, and would then no longer respond to such apparent commitment by playing L. This opens the door for a new type of player 1 to appear, emitting a new signal that is reliably associated with a commitment to B and hence inducing L from player 2. But then the incentive to produce a new mimic appears, and on we go. It appears as if the result could well be a never-ending cycle, as in Robson [117].

In our view, the indirect evolutionary approach will remain incomplete until the evolution of preferences, the evolution of signals about preferences, and the evolution of reactions to these signals, are all analyzed within the model. Perhaps there are outcomes in which players can effectively make commitments by exhibiting the appropriate observable preferences, and there is some force barring the evolutionary pressure to produce mimics, giving us a stationary outcome featuring effective commitment. Perhaps instead the outcome is the sort of cyclical arms race envisioned by Robson [117], with our current situation being a point along this cycle in which some aspects of preferences are at least partially observable. The implications of these scenarios could well be quite different. Further work is required before we have a good idea of what these implications might be. Given the presence of mimics in the natural world, the topic is clearly important. But without more work along these lines, we regard the indirect evolutionary approach as incomplete.

---

[16] For introductions see Harper [69] and Maynard Smith [91, pp. 85–87].

3 What Sort of Preferences?

A representation of preferences such as (1) combines a number of different features, including the choice of what to include as the arguments of the utility function, attitudes toward risk, and trade-offs between consumption at different times. We find it most convenient to address these features separately. We begin in this section by taking it for granted that we can reasonably think of preferences as being defined over a single homogeneous consumption good. We then break our investigation into two parts.

First, we strip away intertemporal considerations to focus on preferences over consumption within a single period. What form do we expect the function $u(c)$ to take? What attitudes toward risk might have evolved? How might risk attitudes vary with one's circumstances or characteristics?

Second, we examine preferences over intertemporal tradeoffs. How do we expect preferences to be aggregated over time? Should we expect preferences to be reasonably approximated by an additively separable utility function, as in (1)? If so, should we expect people to discount the future exponentially? At what rate? If not, how might we expect their discounting to depart from exponential? These questions are all the more pertinent in light of the recent explosion of interest in behavioral economics, much of which is built on the presumption that agents do *not* discount exponentially (cf. Frederick, Loewenstein and O'Donoghue [49]).[17]

3.1 Risk

3.1.1 Attitudes Toward Risk

The expected utility theorem has pride of place in the economic theory of behavior under risk. Whether one believes that expected utility maximization faithfully describes behavior or not, its salience in economic analysis is inescapable.

At first blush, it seems that evolution would surely induce preferences that can be characterized by expected utility maximization.[18] To focus on choice under risk, let us consider a setting in which agents have to choose a lottery from a set of possible lotteries, with the outcome of their selected

---

[17] See Ainslie [1], Loewenstein and Prelec [86], and Loewenstein and Thaler [88] for treatments of present-biased preferences. See Rubinstein [130] for an alternative perspective. Early studies of present bias and self control by Pollak [108], Schelling [138], and Strotz [149] have engendered a large literature. For a few examples, see Elster [39], O'Donoghue and Rabin [100, 101], and Thaler and Shefrin [152].

[18] This section draws on Robson [119].

lottery determining the number of their offspring. The lottery choice is the behavior that is shaped by evolution, being a heritable feature that is passed on from one generation to the next. We then think of a population made up of a number of different types of people, with each type characterized by their choice of economic lottery. All risk is independent across types and individuals, a case that we refer to as "idiosyncratic" (as opposed to "aggregate") risk. For simplicity, we adopt the common assumption that all reproduction is asexual, or "parthenogenetic."[19]

Lotteries are defined over a set of allocations C. The bundle $c \in C$ produces the same expected offspring $\Psi(c)$, regardless of the type of agent, i.e., regardless of the lottery from which this bundle was drawn. Hence, ex ante menus have no ex post consequences. Let q_k^i be the probability that the lottery chosen by type i produces the outcome c_k^i. It follows that the expected number of offspring of type i is then

$$\sum_k q_k^i \Psi(c_k^i).$$

Since the population is large and all risk is idiosyncratic, this is also the growth rate of type i. Thus the most successful type will be the type that maximizes this criterion. But this is simply the maximization of expected utility, where the role of the von Neumann-Morgenstern utility function u is played by the biological production function Ψ.

This evolutionary foundation for expected utility maximization is critically dependent on all the risk being idiosyncratic, or independent across individuals. There seems no compelling reason why all risk should be idiosyncratic. One often begins with hunter-gatherers when thinking about evolution, in an effort to imagine the circumstances under which much of the evolution affecting our current behavior has occurred. Some of the risk in a hunter-gatherer society undoubtedly concerned the weather, which clearly is a shared form of risk. This remained a source of correlated risk as people made the shift to agriculture, perhaps becoming all the more important in the process. In a modern setting, there continue to be important shared risks. Aggregate shocks in the weather have escalated to the possibility of global climate change sufficiently serious as to threaten our survival, while

---

[19] We emphasize that are not under the illusion that human reproduction is asexual, nor do we believe that one can consistently ignore the sexual nature of reproduction when studying evolution. However, models of sexual reproduction are significantly more complicated, and doing justice to sexual reproduction often leaves little analytic headroom to consider other issues. It is thus common practice to more effectively focus on an issue of interest by working with asexual reproduction.

recent events have made it all too clear that social institutions such as financial markets give rise to new sources of correlated risks.

Intuitively, idiosyncratic risk corresponds to having a separate, personal coin flipped for each individual in each period. To keep things simple, let us assume that aggregate risk gives us the opposite extreme in which a single public coin is flipped in each period—heads everyone wins, tails everyone loses. What difference would this make?

To answer this question, let us warm up by considering a related puzzle. An investor much choose between three alternatives:

(1) Investment 1 pays $(3/2)^{52} \simeq \$1,400,000,000$;

(2) Investment 2 pays the expected value of the following lottery. One begins with a balance of one dollar. One then goes through a well-shuffled deck of cards, with 26 black and 26 red cards, successively turning over each card. Each time a red card turns up, the current balance is doubled, while each time a black card comes up, there is no change in the running total;

(3) Investment 3 matches Investment 2, except that the 52 draws are taken from an infinite deck of cards, half red and half black, much like the decks used by casinos to thwart card counters at the blackjack table.

The expected value of Investment 1 is trivially $(3/2)^{52}$, since there is no randomness here. What is the expected return from turning over the first card in Investment 2? 3/2. After that, things get more complex, because it depends now on whether the first draw was red or black. But surely it can't be too bad to take the Investment 2? Surely the expected value of the Investment 2 is something close to $\$1,400,000,000$, even if this is not the exact value?

Compared to the first alternative, Investment 2 is terrible. Indeed, the "lottery" defining Investment 2 involves no uncertainty at all. The payoff is exactly $2^{26} = (\sqrt{2})^{52} \simeq \$67,000,000$, because there are 26 red cards and the doubling effect of each red card is independent of where it arises in the deck. A priori, each card in the deck is equally likely to be red or black, so that the first draw generates an expected value of 3/2. However, the subsequent draws are not independent across cards, and this dependence matters.

Now consider Investment 3. This investment really is a lottery, with realizations that are independent across cards. It no longer matters to subsequent draws whether the first draw is red or black, since there is an infinite number of each color. It is not hard to show that the expected value of the lottery after 52 draws is $(3/2)^{52}$, matching that of the first alternative. To a risk-neutral investor, the two options are then precisely equivalent. A risk-averse investor would choose the first alternative in order to avoid the risk inherent in the third.

Nothing that is fundamental in these comparisons depends upon there being only 52 cards, with a similar comparison holding for any finite number T of draws. The lesson to be learned from this example is that when computing the effect of a series of random variables that accumulate multiplicatively, correlation matters. Notice that if instead the investments were additive—the first adding 3/2 to the running total in each period, and the second being equally likely to add 0 or add 2—then correlation would be irrelevant. The expected payoff of both alternatives would be $(3/2)T$. Indeed, the correlation induced by the 52-card deck, by eliminating any randomness from the problem, would make the two alternatives identical. The infinite deck would preserve the expected value, but make the third alternative riskier.

Now let us turn to an evolutionary setting where analogous forces will appear. We consider a population consisting of two types of infinitely-lived individuals, who differ in the lotteries that govern their number of offspring. In each period, type 1 has either 2 offspring, an event that occurs with probability 1/2, or has only a single offspring, also with probability 1/2. Importantly, all of the risk here is idiosyncratic, meaning that it is independent across all individuals and dates. Type 2 similarly has either 1 or 2 offspring, with each alternative occurring with probability 1/2. However, the risk is now aggregate—either all the type 2 individuals alive at a particular date have two offspring, or they all have only a single offspring—though it remains independent across dates.

One's first reaction here might well be that there should be no difference in the evolutionary success of the two types. From an individual's point of view, the various lotteries involved in each type are identical, making one or two offspring equally likely in each period, independently of how many offspring have appeared in previous period or are expected to appear in subsequent periods. Nonetheless, the two types of individuals face decidedly different evolutionary prospects.

If the population is sufficiently large, then with very high probability, the population ends each period with half again as many type 1s as it began. Because the offspring lotteries are independent across periods, this is an immediate implication of the law of large numbers. Hence, the number of type 1s grows essentially deterministically by a factor of 3/2 in every period, with the number of type 1s at date T being arbitrarily close to $N(T) =$

$(3/2)^T$ (normalizing $N(0)$ to equal 1). The corresponding continuously-compounded growth rate is $\frac{1}{T}\ln N(T) = \ln(3/2)$. The type-1 individuals are thus essentially facing the first alternative in our investment problem.

The number of type 2s is inescapably random, even when the population is extraordinarily large, since in each period a single flip of the offspring coin governs the outcome for every individual. These draws are independent over time, so type 2s are facing the third investment option, played with an infinite deck. It is then not hard to calculate the expected type-2 population size $\tilde{N}(T)$ at time T, finding that $E(\tilde{N}(T)) = (3/2)^T$. This matches the expression for type 1, confirming that the expected number of descendants under each scheme are the same. However, type 2s face risk, with the realized number of type 2s being $\tilde{N}(T) = 2^{\tilde{n}(T)}$, where $\tilde{N}(0) = 1$ and $\tilde{n}(T)$ is the random variable describing the number of heads in a sequence of T flips of a fair coin.

What is the effect of this risk? We can calculate a continuous, deterministic growth rate that reliably describes the behavior of the population as T gets large. In particular, $\frac{1}{T}\ln \tilde{N}(T) = \frac{1}{T}\tilde{n}(T)\ln 2 \to \frac{1}{2}\ln 2 = \ln\sqrt{2}$, with probability one, as $T \to \infty$ (again, by the strong law of large numbers). Hence, while the expected number of type 2s matches the expected number of type 1s, with arbitrarily high probability the realized number of type 2s performs as in Investment 2. Of course, $\sqrt{2} < 3/2$ which implies that with probability one, the ratio of type-1 to type-2 agents goes to infinity. In a strong sense, then, the first type outperforms the second.

What lies behind this comparison? The correlation in the outcomes of Investment 2, whereby every red card calls forth a compensating black card, forces its payoff below that of Investment 1. The independent draws of Investment 3 break this correlation, but over long periods of time the numbers of red and black cards are nonetheless very nearly equal. On outcomes where this is the case, the payoff of Investment 3 falls below that of Investment 1, and similarly the numbers of type 2s fall behind those of type 1s. Investment 3 achieves an expected payoff matching that of Investment 1 by riskily attaching extraordinarily large returns to extraordinarily unlikely events (involving preponderances of red cards). From an evolutionary point of view, this strategy is deadly. With probability arbitrarily close to 1 (for large T), type 2s become a vanishingly small proportion of the population, despite the fact that the expected values of the two are precisely the same. Indeed, with probability one the mean number of type-2 agents grows faster than does the number of type-2 agents itself!

An early use of the word "martingale" was to describe the following betting strategy, mentioned by Casanova in his memoirs: Bet \$1 on a fair coin (or 1 sequin in Casanova's memoirs).[20] If you win, quit, in the process having gained \$1. If you lose, bet \$2 on the next throw. If you win, quit, having gained \$2 − \$1 = \$1. If you lose, bet \$4 on the next throw, and so on. This strategy is claimed to ensure you win \$1.[21]

The martingale betting strategy shares some features with our erstwhile type 2s. Consider the possible outcomes of the martingale strategy after a maximum of $T+1$ flips of the fair coin. One possibility is that you have lost every flip. That is, you might have lost $1+2+...+2^T = 2^{T+1}-1$.[22] The probability of this loss is the probability of $T+1$ heads, or $\left(\frac{1}{2}\right)^{T+1}$. The only other possibility is that you have won, possibly stopping at some earlier time S. If you win, the amount won is always $1 = 2^S - (1+...+2^{S-1})$. The probability of winning must be $1-\left(\frac{1}{2}\right)^{T+1}$. The expected change in wealth is $-\left(\frac{1}{2}\right)^{T+1}\left(2^{T+1}-1\right)+1-\left(\frac{1}{2}\right)^{T+1}=0$, as one would expect— you can't string together a finite series of finite fair bets, no matter how you do it, and expect to do any better than breaking even.[23]

In the limit as $T \to \infty$, however, this is no longer true. The probability of losing tends to zero and that of winning tends to one. In the limiting distribution to which this process converges, you win \$1 for sure. Thus, the limit of the means, \$0, is not equal to the mean of the limiting distribution, \$1. How can this happen? The distribution after a finite number of flips puts a very small probability weight on a very large loss. This yields a non-

---

[20] A sequin was a small gold coin used in Italy. Its value became debased over time, and the word entered English with its current meaning of a dress ornament.

[21] Casanova initially did well with this system, writing that "Before leaving, M– M– asked me to go to her casino, to take some money and to play, taking her for my partner. I did so. I took all the gold I found, and playing the martingale, and doubling my stakes continuously, I won every day during the remainder of the carnival. I was fortunate enough never to lose the sixth card, and, if I had lost it, I should have been without money to play, for I had two thousand sequins on that card. I congratulated myself upon having increased the treasure of my dear mistress, who wrote to me that, for the sake of civility, we ought to have a supper 'en partie carrée' on Shrove Monday. I consented." (This quotation is from Chapter 21 of *The Complete Memoirs of Jacques Casanova de Seingalt, Volume Two: To Paris and Prison*, translated by Arthur Machen, published by G. P. Putnam's Sons of New York, and available at http://www.gutenberg.org/files/2981/2981-h/v2.htm.)

[22] To confirm this expression, suppose it holds after losing T times. It follows that it holds after losing $T+1$ times because $1+2+...+2^{T+1} = 2(2^{T+1}-1) = 2^{T+2}-1$.

[23] It seems that Casanova came to a similar conclusion, writing in Chapter 24 that, "I still played on the martingale, but with such bad luck that I was soon left without a sequin. As I shared my property with M– M– I was obliged to tell her of my losses, and it was at her request that I sold all her diamonds, losing what I got for them; she had now only five hundred sequins by her. There was no more talk of her escaping from the convent, for we had nothing to live on! I still gamed, but for small stakes, waiting for the slow return of good luck."

vanishing contribution to the mean. In the limit, however, the probability of this loss converges to zero, giving us an upward jump in the mean "at the limit."

In our simple biological example, the mean of the type 2 population is similarly (if inversely) held *up* by very small probabilities of very large populations. In the limit, these probabilities vanish, so the growth of the population is overestimated by the mean. Despite having the same mean, the population almost surely fares worse under aggregate uncertainty (the type 2s) than under individual uncertainty (type 1).

The implication of this difference is that evolutionarily optimal strategies should be be more averse to aggregate risk than to equivalent idiosyncratic risk, in the sense that people should be less willing to accept lotteries incorporating aggregate risks. From an individual point of view, this may seem bizarre. Why should I be on the verge of undertaking an investment, only to balk upon learning that my realizations will be shared by many other people? But we can expect evolution to have learned via experience that such investments are to be shunned, and can expect this to be reflected in our preferences.

The example can be recast as an economic choice as follows. Suppose that bundles c_1 and c_2 induce the offspring levels 1 and 2, so $\Psi(c_1) = 1$ and $\Psi(c_2) = 2$, where Ψ is the common production function for expected offspring. Now individuals must choose between lottery 1 and lottery 2. Lottery 1 yields c_1 and c_2 each with probability 1/2, where all this risk is independent. Lottery 2 also yields c_1 and c_2 each with probability 1/2, but now all this risk is aggregate. From an expected utility point of view, these two lotteries should be equivalent. Indeed, even from the perspective of any decision theory that applies the apparently weak notion of "probabilistic sophistication," these two lotteries should be equivalent. But it is not enough here to consider only one's own payoffs and the associated probabilities, as such sophistication requires. One must also consider how the uncertainty affects others. That is, preferences are interdependent. In an evolutionary environment, individuals should prefer lottery 1 to lottery 2.

The most general case that can easily be analyzed is as follows. Given an aggregate environment z, each type i faces an idiosyncratic economic lottery where $q_k^{i,z}$ is the probability of receiving a commodity bundle $c_k^{i,z}$. We let $\Psi(c)$ be the expected offspring from bundle c for any state and any type, where any underlying risk here is also idiosyncratic. Hence $\sum_k q_k^{i,z}\Psi(c_k^{i,z})$ is the expected offspring of type i in state z. If each state z has probability ρ_z, then the long run limiting exponential growth rate of type i is

$$\sum_z \rho_z \ln\left(\sum_k q_k^{i,z}\Psi(c_k^{i,z})\right). \tag{3}$$

Hence the type that maximizes this expression should be favored by natural selection. In particular, we see the preference for idiosyncratic rather than aggregate risk in our example, since

$$\ln\left((1/2)\Psi(c_1) + (1/2)\Psi(c_2)\right) > (1/2)\ln\Psi(c_1) + (1/2)\ln\Psi(c_2),$$

by the strict concavity of the function ln.

What are the behavioral implications of the distinction between aggregate and idiosyncratic risk? People may strictly prefer to take idiosyncratic lotteries for reasons that are quite distinct from a conventional explanation in terms of the convexity of the von Neumann-Morgenstern utility. Perhaps the simplest example of this is due to Cooper and Kaplan [27]. Consider the evolutionary success of a parthenogenetic animal. Suppose the probability of a snowy winter is $\rho \in (0, 1/2)$ and hence the probability of a clear winter is $1-\rho \in (1/2, 1)$. The animal is hunted by predators whom it hopes to escape by blending indistinguishably into its surroundings. As a result, animals with dark coats survive clear winters but die in snowy winters, while those that develop white coats survive snowy winters but die in clear ones. Clearly a type that always has a dark coat is doomed to extinction with the first white winter, and one that always has a white coat is doomed by the first clear winter. Suppose the chameleon-like strategy of changing colors with the nature of the winter is infeasible. Then consider a type whose members randomize—choosing a white coat with probability π and a dark coat with probability $1-\pi$. That is, all individuals of this type are genetically identical, where this means merely that they choose their winter color from the same idiosyncratic lottery, but experience different ex post outcomes. The overall growth rate of this type is then

$$r = \rho\ln\pi + (1-\rho)\ln(1-\pi),$$

which is readily shown to be maximized by choosing $\pi = \rho$. In particular, such "probability matching" allows this type to avoid extinction.

This argument is developed further by Bergstrom [13], who casts the story in terms of squirrels who might similarly adopt a mixed strategy in saving food for a winter of variable length. Even if the long and harsh winters are extraordinarily rare, a pure type that stored enough food only for shorter and milder winters would be doomed to extinction, while a pure strategy of saving for the longest and harshest of winters is very wasteful, consuming resources and incurring risks to accumulate food that virtually always goes unused. The optimal response is a mixture in which only a small fraction of the population stockpiles sufficient food to ensure the worst of winters, allowing the population to avoid extinction while most members also avoid overwhelmingly wasteful accumulation.

Cooper and Kaplan [27] interestingly interpreted individuals who choose a white coat in their model after the flip of their coin as being "altruistic." Why? The probability of such an individual dying in their model is higher than the probability of death for an individual with a dark coat, simply because $1 - \rho > 1/2 > \rho$. The apparent altruism thus arises out of a choice that seems to decrease an agent's probability of survival, while protecting the population from extinction. Why would such an agent ever make a choice? Why not maximize the probability of survival? Before we can interpret this choice as altruism, we must make sure of the correct notion of fitness (as a biologist would put it) or, equivalently, the correct utility function.

Grafen [63] offers a resolution of the apparent altruism puzzle raised by Cooper and Kaplan. Consider a continuum of agents of size 1. Suppose π of these agents choose white and $1 - \pi$ choose dark. Now consider the choice of a small mass of individuals of size ε. If they choose white, the *expected fraction* of the population they will constitute at the end of the winter is $\frac{\rho\varepsilon}{\pi}$, which equals ε if $\rho = \pi$. If they choose dark, the *expected fraction* of the population they will constitute is $\frac{1-\rho}{1-\pi}\varepsilon$, which again equals ε if $\rho = \pi$. Each individual of the type that randomizes $(\rho, 1-\rho)$ thus maximizes the expected fraction of the population it will comprise, and this expected fraction of the population is the appropriate notion of biological fitness. Death brings zero fitness no matter what the state of the population, but when you survive it matters how large you loom in the population.

To reinterpret this from an economic point of view, the result is that the usual selfish preferences are inadequate in explaining behavior in the face of aggregate uncertainty. It is instead important to consider not only the likelihood of death, but also how well you are doing when you do survive *relative to others*. The the appropriate notion of utility must then be interdependent. See Curry [30] for an analysis of this interdependence.

3.1.2 Risk and Status

It is a common observation that people exhibit risk-aversion when making some choices while also exhibiting risk-preference in other cases. People buy both insurance and lottery tickets. The standard explanation for this behavior begins with Friedman and Savage [54], who suggested that the typical von Neumann-Morgenstern utility function is concave over low values of wealth but then becomes convex over higher values. People with such utility functions would seek insurance protection against downside risk, while at the same time buying lottery tickets that promise a small probability of a large increase in wealth. One can account for the observation that actual lotteries have a nontrivial array of prizes, rather than a single grand prize, by assuming that there is a final range of wealth over which von Neumann-Morgenstern utility is again concave.

The Friedman-Savage explanation views utility as being defined over absolute wealth levels. The difficulty here is that absolute wealth levels have changed dramatically over a relatively short period of our recent history. If a Friedman-Savage utility function supported the simultaneous purchase of insurance and gambling in a particular society at a particular date, then growing wealth levels would make it difficult to use the same utility function in explaining similar phenomena at a later date. Indeed, if utility functions are stable, then the market for insurance should wither away, as the number of individuals in the requisite low range of wealth decrease. Lotteries may also have diminishing prizes over time, since a lower prize would attain the same target level of final wealth. Nothing in our current experience suggests that the demand for insurance has dissipated as our society has gotten wealthier, or that lottery prizes are deteriorating.

The preceding argument relies on a particularly simple utility function, and one could come closer to a consistent model of behavior with a more elaborate function. In the process, of course, one must worry about constructing ever-more-sophisticated models that ultimately collapse under the weight of their complexity, just as epicycles ultimately gave way to a more parsimonious model. A seemingly more likely explanation is that utility functions have changed over time. Increasing wealth has not vitiated the need for insurance because utility functions have ratcheted up along with wealth levels. While intuitive, this explanation alone is discomforting in its reliance on the exogenously generated shifting of utility functions. Why do our utility functions change as our society gets wealthier? When is this shift likely to be especially pronounced, and when is it likely to be attenuated? What implications does it have for behavior, and for economic policy?

Robson [118] (see also Robson [120]) offers a model that allows us to address these types of questions. The key ingredient is that people care not only about their absolute wealth, but also about their position in the wealth distribution.[24] There are many reasons why people might care about how their wealth compares to that of others. For the purposes of this discussion, we simply assume that people care about "status," which in turn is determined by their place in the wealth distribution. We close this section with some examples of the considerations that might give rise to such a concern for status, deferring a more careful discussion to Section 4.2.

We suppose that an individual with wealth w attains status $S = F(w)$, where F is the continuous cumulative distribution function describing the wealth distribution in the relevant population. The population is represented by a continuum, normalized to have size 1. Hence status is the proportion of individuals that the individual outranks in terms of wealth. The individual has a von Neumann-Morgenstern utility function that is concave in w but convex in S. The convexity of S, indicating that increases in status are especially valuable near the upper end of the wealth distribution, will lead to risk-seeking behavior over some wealth levels.

For convenience, let us work with a particular functional form, given by:

$$u(w, S) = \ln w + kS^{\beta},$$

where $k > 0$ and $\beta \geq 2$. Suppose, for simplicity, that the wealth distribution is uniform on the interval of wealth levels $[0, \gamma]$, and hence is given by

$$\begin{array}{rrcl} & F(w) & = & w/\gamma \text{ for all } w \in [0, \gamma] \\ \text{and} & F(w) & = & 1 \text{ for all } w > \gamma. \end{array}$$

In a more complete model, of course, one would want the distribution of wealth levels to be endogenous, but a partial-equilibrium approach will serve us well here.

Suppose now that we condense the utility function so that it takes only wealth as an argument by defining $v(w) = u(w, F(w))$. Then it follows that

$$\begin{array}{rcl} v''(w) & < & 0 \text{ for all } w \in (0, \tilde{w}), \quad \text{where } \tilde{w} = \dfrac{\gamma}{(\beta(\beta-1)k)^{1/\beta}} \\ v''(\tilde{w}) & = & 0 \\ v''(w) & > & 0 \text{ for all } w \in (\tilde{w}, \gamma) \text{ and} \\ v''(w) & < & 0 \text{ for all } w > \gamma, \end{array}$$

---

[24] A similar convention is helpful in accounting for patterns of consumption as a function of wealth or income, as was pointed out long ago by Duesenberry [34]. See Rabin [111] and and Cox and Sadiraj [28] for another discussion of whether utility is usefully defined over absolute wealth levels.

where we assume that $\beta(\beta-1)k > 1$ so that $\tilde{w} < \gamma$.

This example yields the concave-convex-concave utility described by Friedman and Savage. The convexity of $u(w, S)$ in S is needed to obtain the intermediate range of wealth, $(\tilde{w}, \gamma)$, over which $v(w)$ is convex. The concavity of $u(w, S)$ in w yields the concavity of $v(w)$ over the initial and final ranges $(0, \tilde{w})$ and (γ, ∞). The latter range appears despite the status effect because $f(w) = 0$ on (γ, ∞). Note that the first inflection point, $\tilde{w}$, can fall anywhere in $(0, \gamma]$, depending on the values of the parameters.

This model allows us to capture behavior that is risk-averse over some income ranges and risk-seeking over others, without such counterfactual implications as the prediction that the insurance industry will wither away as a society become wealthier. Consider, for example, a uniform multiplicative shift in the wealth distribution, represented by an increase in γ. The inflection point $\tilde{w}$ is subject to the same multiplicative shift, so the same individual lies on the watershed between risk-aversion and risk-preference. Similarly, this model is consistent with prizes in lotteries that grow over time in step with the growth of the wealth distribution. That is, the wealth level γ marking the transition from risk-preference to risk-aversion is subject to this same shift.[25] To an analyst using models based on utility functions of the form $v(w)$ to study the economy, it would look as if the parameters of the utility functions are adjusting at about the same rate as wealth is growing, in the process coincidentally preserving the qualitative features of behavior. In fact, however, there would be nothing exogenous in the seemingly shifting utilities.

If the von Neumann-Morgenstern utility of wealth alone has a concave-convex-concave shape, as in Friedman and Savage, and individuals have access to a variety of fair bets, then individuals in an intermediate range will find it attractive to take gambles whose outcomes will put them either into a low initial range of wealth or a high terminal range (e.g., Friedman [53]). As a result, the middle class should disappear. However, Robson [118] shows that if the von Neumann-Morgenstern utility also depends on status, this redistribution of wealth will end before the middle class is completely depopulated. Robson [118] also discusses how a concern with status in this sense involves an externality. If we contemplate the effects of an increase in our wealth, we take into account the effect this has in increasing our status, but we neglect the effect it has in lowering other individuals' status. There

---

[25] This argument can be immediately generalized to utility functions of the form $u(w, S) = \ln w + v(S)$, where v is any increasing differentiable function and to an arbitrary continuous cumulative distribution function of wealth F.

may well then be too much gambling. Less obviously, there may instead be too little—there are distributions of wealth that are stable, in the sense that no one wishes to take any fair bet, despite the existence of fair bets that induce a Pareto improvement.

How might the concern with status that lies at the heart of this model have evolved? We only sample the many possibilities here. For example, Robson [120] considers how a concern for status and an attendant risk-preference might arise in a polygynous setting, where females choose males based on their wealth. Cole, Mailath and Postlewaite [26] suggest that concerns for status may arise because some goods in our economy are allocated not by prices, but by nonmarket mechanisms in which status plays a role. Cole, Mailath and Postlewaite suggest the "marriage market" as a prime such example, where access to desirable mates often hinges on placing well in a status ordering that depends importantly on wealth. Additional points of entry into the large literature include Becker, Murphy and Werning [10], Frank [47], and Ray and Robson [113].

What form might a concern with status have? There are two intriguing possibilities. If access to desirable mates lies behind a concern for status, then evolution may have designed us with utility functions that depend directly only on absolute wealth and mates. The contest for mates may give rise to behavior that makes it look as if people have a concern for relative wealth, but this concern would be instrumental rather than intrinsic (cf. Postlewaite [109]). Hence, status may be important, while the standard economists' inclination to work with "selfish" preferences, or preferences only over one's own outcomes may still have a solid biological foundation. Alternatively, constraints on the evolutionary design process, perhaps rising out of information or complexity considerations, may cause evolution to find it easier or more expeditious to simply design us with preferences over relative wealth, trusting that this will lead (perhaps more reliably) to the appropriate outcomes. In this case the concern with relative wealth is intrinsic and we are pushed away from the familiar selfish preferences.

Determining which aspects of our preferences are instrumental and which are intrinsic is an important and challenging question. We return to the possibility that status may play a role in preferences in Section 4.2

3.1.3 Implications

Where do we look for the implications of these evolutionary models, implications that Section 2.2 suggested should be the signature of the evolutionary approach? One obvious point stands out here. People should evaluate idiosyncratic and aggregate risks differently.

A standard finding in psychological studies of risk attitudes is that a feeling of control is important in inducing people to be comfortable with risk.[26] Risks arising out of situations in which people feel themselves unable to affect the outcome cause considerably more apprehension than risks arising out of circumstances people perceive themselves to control. People who fear flying think nothing about undertaking a much more dangerous drive home from the airport.[27] The risk of a meteor strike that eliminates human life on Earth is considered more serious than many other risks with comparable individual death probabilities. Why might this be the case? The first task facing evolution in an attempt to induce different behavior in the face of idiosyncratic and aggregate risks is to give us a way of recognizing these risks. "Control" may be a convenient stand-in for an idiosyncratic risk. If so, then our seemingly irrational fear of uncontrolled risk may be a mechanism inducing an evolutionarily rational fear of aggregate risk.

3.2 Time

We now turn our attention from the within-period considerations captured by the function $u(c)$ to the question of intertemporal trade-offs. In doing so, we strip away all considerations of the nature of $u(c)$ by focussing on preferences over offspring. Hence, the agents in our model will do nothing other than be born, have offspring, and then die. In addition, no notion of the quality of offspring will enter our discussion. Agents will differ only in the number and timing of their offspring.

Our motivation in constructing such a model is to work with as close a link as possible between the model and the criteria for evolutionary success. The ultimate goal of evolution is successful reproduction. As simple as this sounds, "reproduction" is a multifaceted process and "success" involves managing a variety of tradeoffs. We eliminate many of these tradeoffs by working with a world of homogeneous offspring, focussing attention on the twin objectives of having many offspring and having them quickly. How does evolution balance "many" versus "quickly?" We view this as the obvious place to look for clues to how our preferences treat intertemporal tradeoffs, and so this becomes the focus of our analysis.

---

[26] See Slovic, Fischhoff and Lichtenstein [144] for an early contribution to this literature and Slovic [143] for a more recent introduction.

[27] Indeed, Gigerenzer [58, pp. 31] suggests that direct death toll in the September 11, 2001 attack on New York's World Trade Center may have been surpassed by the increased traffic deaths caused by subsequent substitution of driving for air travel.

Evolution must not only identify the preferred mix of number and timing of offspring, but also solve the problem of how to induce this behavior. As faulty as it is, introspection suggests that evolution has not accomplished her goal by having us make constant calculations as to whether our next restaurant choice will increase or decrease the number of children we expect, or whether our choice of what car to drive will advance or postpone our next child. Instead, evolution works through utility functions that attach rewards to a host of intermediate goals, such as being well nourished. How and why evolution has done this is again an important and fascinating question, but is swept out of sight here.

Our basic notion is that of a "life history." A life history specifies the number of offspring born to an agent at each of the agent's ages. We assume that such life histories are heritable. The evolutionary approach proceeds by asking which life history will come to dominate a population in which a variety of life histories are initially present. In particular, we imagine mutations regularly inserting different life histories into a population. Some cause the group of agents characterized by such a life history to grow rapidly, some lead to slow rates of growth. The life history leading to the largest growth rate will eventually dominate the population. Having found such a life history, we will be especially interested in characterizing the implicit intertemporal trade-offs.

The question of why people discount is an old one. It seems intuitively obvious that future consumption is less valuable than current consumption, but why is this the case? A good place to start in one's search for an answer is the work of Fisher [43, pp. 84–85], who pointed to one reason future rewards might be discounted—an intervening death might prevent an agent from enjoying the reward. This gives us a link between mortality and discounting that has often reappeared (e.g., Yaari [166]), and that will again arise in our model. Hansson and Stuart [68] and Rogers [129] (see also Robson and Szentes [127]) point to a second factor affecting discounting. They construct models in which evolution selects in favor of people whose discounting reflects the growth rate of the population with whom they are competing. Our first order of business, in Section 3.2.1, is to put these ideas together in the simplest model possible, leading to the conclusion that evolution will induce people to discount exponentially at the sum of the population growth rate and mortality rate. We then consider a sequence of variations on this model.

3.2.1 A Simple Beginning: Semelparous Life Histories

We begin by considering only *semelparous* life histories, in which an organism reproduces at a fixed, single age (if it survives that long) and then dies.[28] We do not view this model as a realistic foundation for understanding discounting, but it does provide a revealing introduction to the relevant evolutionary forces.

A life history in this context is simply a pair (x, τ), where x is the agent's expected number of offspring and τ is the age at which these offspring are produced. The agents in this environment live a particularly simple life. They wait until age τ, possibly dying beforehand, and then have x offspring. At that point the parents may die or may live longer, but in the latter case do so without further reproduction. We need not choose between these alternatives because the possibility of such a continued but barren life is irrelevant from an evolutionary point of view. Agents who survive past their reproductive age may increase the size of the population at any given time, but will have no effect on the population growth rate. As a result, any mutation that sacrifices post-reproduction survival in order to increase the number of offspring x or decrease the age τ at which they are produced will be evolutionarily favored, no matter what the terms of the trade-off.

In the parlance of evolutionary biology, the particularly simple life histories of these agents earns them the title of "Darwinian dolts" (cf. Stearns and Hoekstra [148, p. 219]). In particular, if reproduction is affected by aggregate risks, such as predators or plagues that threaten survival to reproductive age, famines that threaten the ability to produce offspring, or climatic fluctuations that threaten offspring survival, then a semelparous life history can expose its practitioners to costly risk. Nonetheless, there is much to be learned from Darwinian dolts.

We examine a group of agents whose members are all characterized by a particular life history (x, τ). We will speak throughout as if a life history is a deterministic relationship, with each age-τ parent having precisely x offspring. The interpretation is that x is the *expected* number of offspring born to age-τ parents. As long as the group size is sufficiently large and the random variables determining the number of offspring born to each parent are independent, then the average number of offspring will be very close to x and x will provide a very good approximation of the behavior of the evolution of the population.[29] The life history (x, τ) is presumably the result

---

[28] This section is based on Robson and Samuelson [124].

[29] For typical limit theorems underlying this type of deterministic approximation, see Benaïm and Weibull [11]. The case of a continuum of agents raises technical problems.

of various choices on the part of the agent, such as where to seek food, what food to eat, when to mate, what sort of precautions to take against enemies, and so on, all of which have an important effect on reproduction, but which do not appear explicitly in our model.

An agent who delays reproduction increases the risk of dying before reaching reproductive age. In particular, an agent choosing (x, τ) survives for the length of time τ required to reach reproductive age with probability $e^{-\delta\tau}$, where δ is the instantaneous death rate. If and only if the agent survives, the x offspring appear.

Consider a population characterized by strategy (x, τ), of initial size N_0. How large will this population be at time $t > 0$? Let us follow a dynasty, meaning a cohort of agents initially of some age τ', who have offspring when they reach age τ, with these offspring then having their offspring upon reaching age τ, and so on. From time 0 until time t, there will have been approximately (depending on the cohort's initial age and integer problems) t/τ intervals during which this dynasty will have first shrunk by factor $e^{-\delta\tau}$, as the population is whittled away by death while awaiting its next opportunity to reproduce, and then multiplied itself by x as it reproduces. The population at time t is thus

$$N_0 \left(e^{-\delta\tau} x \right)^{\frac{t}{\tau}} .$$

The growth factor for this population is then $e^{-\delta}(x)^{\frac{1}{\tau}}$.

If the population is characterized by a variety of life histories, then evolution will select for the value (x, τ) that maximizes $e^{-\delta}(x)^{\frac{1}{\tau}}$ or, equivalently, that maximizes

$$\frac{\ln x}{\tau}. \tag{4}$$

Hence, evolution evaluates births according to the function $\ln(\cdot)$ and discounts them hyperbolically. The equilibrium population will grow exponentially at the growth rate $-\delta + \frac{\ln x}{\tau}$.

Have we just discovered an evolutionary foundation for the hyperbolic discounting that lies at the core of much of behavioral economics? Caution is in order on several counts. First, the phrase "hyperbolic discounting" is used to denote a variety of discounting patterns, many of which do not match (4). Perhaps the most common of these is the "$\beta - \delta$" formulation, in which payoffs in period t are discounted to the present (period 0) at rate $\beta\delta^{t-1}$, with $\beta > \delta$. As a result, the delay between the current and next periods

---

See Al-Najjar [3] for a discussion.

is weighted especially heavily, with subsequent delays being equivalent. In contrast, the preferences given by (4) represent hyperbolic discounting in the literal sense, in that period-t payoffs are discounted to the present by the factor $1/t$. This discounting pattern is common in biological models of foraging (e.g., Houston and McNamara [75, chapter 4], Kacelnik [76], Bulmer [19, chapter 6]), but less common in economics. Second, hyperbolic discounting is especially intriguing to behavioral economists for its ability to generate preference reversals. In contrast, no incentive for preference reversals arises in the present evolutionary context. Indeed, we have not yet built a rich enough set of choices into the model to talk about preference reversals. We have simply identified the criterion for finding the optimal tradeoff between the delay to reproduction and the number of attendant offspring.

More importantly, we need to think carefully about making the leap from (4) to individual preferences. The preferences captured by (4) are relevant for asking a number of questions about the comparative statics of evolution. For example, these preferences are the appropriate guide if we want to know which of two populations, characterized by different life histories, will grow faster, or which of two mutants will be most successful in invading a population. Suppose, however, that we are interested in using preferences to describe the choices we see in a particular population. Let (x, τ) be the equilibrium life history, giving rise to a population that grows exponentially at rate $r = \ln\left(e^{-\delta} x^{\frac{1}{\tau}}\right) = -\delta + \frac{1}{\tau}\ln x$. Then consider the alternative strategy $(\tilde{x}, \tilde{\tau})$. Suppose this alternative strategy is feasible but not chosen (and hence gives a lower growth rate $\tilde{r}$). What preferences would we infer from this observation? We could assume that preferences are given by (4). However, we could also assume that the agents evaluate births linearly and discount exponentially at rate $-(\delta + r)$, so that (x, τ) is evaluated as $e^{-(\delta+r)\tau}x$. In particular, to confirm that such preferences rationalize the choice of (x, τ), we need only note that[30]

$$\begin{aligned} e^{-(\delta+r)\tau}x > e^{-(\delta+r)\tilde{\tau}}\tilde{x} \quad &\Leftrightarrow \quad e^{-(\delta+r)\tau}x > e^{-r\tilde{\tau}}e^{\tilde{r}\tilde{\tau}}e^{-(\delta+\tilde{r})\tilde{\tau}}\tilde{x} \\ &\Leftrightarrow \quad r > \tilde{r}. \end{aligned}$$

Exponential discounting, at the sum of the death and optimal growth rates, thus characterizes the preferences with which evolution will endow her agents. This representation of preferences is intuitive. There are two costs of delaying reproduction. One of these is simply that death occurs at rate δ. The

---

[30] The second inference follows form the observation that $e^{-(\delta+r)\tau}x = 1 = e^{-(\delta+\tilde{r})\tilde{\tau}}\tilde{x}$.

other is that a given number of offspring will comprise a smaller fraction of a population growing at rate r. The sum of these two rates is the rate at which delaying births causes an agent to fall behind the population.

3.2.2 Extensions

With this basic result in hand, we consider six respects in which this analysis is limited, and hence warrants generalization:

1. Once the optimal strategy has spread throughout the population, the population will grow exponentially at the resulting growth rate. In practice, we do not expect populations to grow without bound, and so a model with some constraints on population size would be more reasonable.
2. We have allowed agents to reproduce only once, while we expect situations to be important in which agents can reproduce more than once.
3. Even if reproduction is the ultimate issue of concern to evolution, all of our experience as well as our economic literature suggests that we have preferences over many other things, commonly lumped together in economic models under the label of consumption.
4. The agents in our model are homogeneous, with every agent facing the same set of choices and making the same optimal choice. How do we incorporate heterogeneity into the model?
5. All of the uncertainty in the model is idiosyncratic, and hence washes out in the analysis of the population. What if there is aggregate uncertainty?
6. One motivation for studying evolutionary foundations for discounting is to glean insights into models of hyperbolic discounting, present bias, and preference reversal. We have found a hint of hyperbolic discounting in preferences that are relevant for evolutionary comparative statics, but none in the induced individual behavior. Does an evolutionary perspective lock us into exponential discounting?

The following sections examine each of these points in turn.

3.2.3 Environmental Capacity

The discount rate in our analysis is tied closely to the population growth rate. A more rapid population growth induces a higher discount rate, while a population that shrinks sufficiently rapidly will induce negative discounting (in which case reproduction is better deferred). If the population growth rate is zero, agents will discount at the death rate δ.

The difficulty here is that we do not expect populations to grow without bound. If nothing else, an exponentially growing population will eventually produce a physical mass of agents too large to fit on the Earth, even neglecting any considerations of whether the planet can sustain them.[31] In some instances, resource constraints may not bind for a long time. One might then argue that an unconstrained model is a reasonable approximation of our evolutionary past, even if not a good guide to our future. However, we must be wary of appealing to the latter type of short-run argument when interpreting a theory whose predictions consist of limiting results. Perhaps more to the point, it seems likely that environmental constraints restricted human growth rates to be near zero throughout much of our evolutionary past.

Nothing in our analysis changes if we modify the death rate δ to reflect environmental constraints on the population size. We can do so while retaining all of the analysis in Section 3.2.1, as long as we interpret the death rate appearing in our model as the steady-state rate that balances population growth and environmental constraints.

In particular, notice that the discount rate in our exponential-discounting representation of preferences, given by

$$\delta + r = \frac{1}{\tau} \ln x,$$

is *independent* of the death rate. If an increasing population size uniformly increases the death rate, the growth rate will exhibit a corresponding decrease, leaving the discount rate unaffected. The discount rate is affected only by the life-history specification (x, τ). In a sense, we have thus turned the views of Fisher [43] and Yaari [166] on their heads. Instead of being a primary reason for discounting, death has nothing to do with the appropriate

---

[31] Pursuing this point into the more fanciful, the space occupied by an exponentially growing population will eventually contain a sphere whose radius expands at a rate exceeding the speed of light, ensuring that we cannot alleviate the problem by travel to other planets (at least under our current understanding of physics). Finding oneself too heavily involved in such arguments is a reliable sign that something is missing from one's model.

discount rate.[32]

3.2.4 Iteroparous Life Histories

We can easily generalize the analysis to *iteroparous* life histories, in which an individual may have offspring at more than one age. Among other advantages, such a life history may allow individuals to diversify some of the (unmodeled, in our analysis) aggregate risks that might make semelparity particularly precarious.

It is convenient here to let time be measured discretely. Let each agent live for T periods, producing x_τ offspring in each period $\tau = 1, \dots, T$. A life history is then a collection $(x_1, x_2, \dots, x_T)$, where some of these entries may be zero.

Our basic tool for keeping track of the population is a *Leslie* matrix (Leslie [81, 82]), given in this case by

$$\begin{bmatrix} e^{-\delta}x_1 & e^{-\delta} & 0 & \dots & 0 & 0 \\ e^{-\delta}x_2 & 0 & e^{-\delta} & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ e^{-\delta}x_{T-1} & 0 & 0 & \dots & 0 & e^{-\delta} \\ e^{-\delta}x_T & 0 & 0 & \dots & 0 & 0 \end{bmatrix}.$$

Each row $\tau = 1, \dots, T$ in this matrix corresponds to the fate of agents of age τ in the population in each period. The first entry in this row indicates that these agents have x_τ offspring, which survive to become the next period's 1-period-olds at rate $e^{-\delta}$. The second term in the row indicates that at rate $e^{-\delta}$, the agents of age τ themselves survive to become one period older.

Letting X be the Leslie matrix, the population at time t is given by

$$N'(t) = N'(0)X^t, \tag{5}$$

where $N'(t)$ is a (transposed) vector $(N_1(t), \dots, N_T(t))$ giving the number of agents in the population of each age $1, \dots, T$ at time t. The fate of the population thus hinges on the properties of X^t. The Perron-Frobenius theorem (Seneta [140, Theorem 1.1]) implies that the Leslie matrix has a "dominant" eigenvalue ϕ that is real, positive, of multiplicity one, and that

---

[32] We must be careful here to distinguish proximate and ultimate causes. The latter are the evolutionary considerations that shape the optimal life history, while the former are the mechanisms by which evolution induces the attendant optimal behavior. The death rate does not appear among the ultimate causes of discounting.

strictly exceeds the modulus of all other eigenvalues.[33] This eigenvalue is the population growth factor, and its log is the corresponding growth rate, in the sense that (Seneta [140, Theorem 1.2])

$$\lim_{t\to\infty} \frac{X^t}{\phi^t} = vu',$$

where the vectors u and v are the strictly positive left ($u'X = \phi u'$) and right ($Xv = \phi v$) eigenvectors associated with ϕ, normalized so that $u'v = 1$ and $\sum_{\tau=1}^{T} u_\tau = 1$.[34]

Evolution must select for behavior that maximizes the eigenvalue ϕ, or equivalently, that maximizes the long-run growth rate $\ln \phi$. This eigenvalue solves the characteristic equation

$$\left| \begin{array}{ccccc} e^{-\delta}x_1 - \phi & e^{-\delta} & 0 & \dots & 0 \\ e^{-\delta}x_2 & -\phi & S & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ e^{-\delta}x_{T-1} & 0 & 0 & \dots & e^{-\delta} \\ e^{-\delta}x_T & 0 & 0 & \dots & -\phi \end{array} \right| = 0,$$

or, equivalently,

$$\Phi = x_1 + \frac{x_2}{\Phi} + \frac{x_3}{\Phi^2} + \ldots + \frac{x_T}{\Phi^{T-1}}, \tag{6}$$

where

$$\Phi = \frac{\phi}{e^{-\delta}}.$$

Equation (6) gives us our basic description of preferences. Evolution will endow an agent with preferences (or more precisely, would endow an agent with behavior consistent with such preferences) whose indifference curves are described by the right side of (6), with Φ corresponding to the optimal growth rate. In particular, choices $(x_1, \ldots, x_T)$ that lead to a smaller value on the right side of (6) would lead to a lower growth rate and would be optimally rejected by the agent.

As with the case of semelparous life histories, we can draw two kinds of conclusions from these results. First, we can ask questions about "evolution's preferences" or "evolutionary comparative statics," addressing the

---

[33] We assume that the Leslie matrix X is primitive, in that there exists some $k > 0$ for which X^k is strictly positive. A sufficient condition for this is that there exist two relatively prime ages τ and τ' for which x_τ and $x_{\tau'}$ are both nonzero.

[34] Regardless of the initial condition $N'(0)$, the proportion of the population of each age τ approaches u_τ. The vector v gives the "reproductive value" of an individual of each age, or the relative contribution that each such individual makes to the long run population.

relative performance of alternative populations or alternative mutants within a population. Here, we once again recover hints of hyperbolic discounting, seen in the fact that the evolutionary criterion for evaluating alternative life histories, given by (6), contains our previous results for semelparous life histories as a special case. In particular, it is immediate from (6) that evolution is indifferent over two semelparous strategies (x_1, τ_1) and (x_2, τ_2) if and only if $x_1^{\frac{1}{\tau_1}} = x_2^{\frac{1}{\tau_2}}$. This confirms that the semelparous analysis is a special case of this more general model. Preferences over the remaining iteroparous strategies are captured by connecting indifferent semelparous strategies with linear indifference surfaces. More generally, this population growth rate is a complex function of the fertility profile. If we let $\Phi = \Phi(x_1, x_2, ...)$ be the function implicitly defined by (6), then the marginal rate of substitution between x_t and x_{t+1} is Φ itself, which is a strictly increasing function of *each* x_τ for $\tau = 1, \ldots, T$. It is then immediate that there can be no additively separable representation of evolution's preferences.

Alternatively, we can ask about the behavior we would observe from agents. Agents can once again be induced to make optimal choices via exponentially discounting offspring at the sum of the death and optimal growth rates. Letting $(x_1, \ldots, x_T)$ be the optimal fertility profile and Φ be implicity defined by (6), we have

$$1 = \frac{x_1}{\Phi} + \frac{x_2}{\Phi^2} + \ldots + \frac{x_T}{\Phi^T}.$$

Now suppose an alternative fertility/utility profile $(x'_1, \ldots, x'_T)$ is feasible but is not chosen because it gives a smaller growth rate. Then

$$\frac{x_1}{\Phi} + \frac{x_2}{\Phi^2} + \ldots + \frac{x_T}{\Phi^T} \quad = \quad 1 \quad > \quad \frac{x'_1}{\Phi} + \frac{x'_2}{\Phi^2} + \ldots + \frac{x'_T}{\Phi^T}.$$

The agent's behavior is thus again consistent with exponentially discounted preferences, with a discount rate given by the sum of the death rate and population growth rate.

3.2.5 Consumption

Economists are typically interested in preferences over consumption rather than births and mortality. Perhaps the simplest way to transform a model of preferences over fertility and mortality rates into a model of preferences over consumption is to assume that births are a function of consumption, so that preferences over consumption are those induced by the underlying preferences over births. Notice that in doing so, we are not assuming that every visit to a restaurant requires a quick calculation as to whether steak or fish is more likely to lead to more offspring. Instead, our presumption is that evolution simply gives the agent preferences over steak and fish, with evolution shaping these preferences to reflect the required calculation.

Consider for simplicity the case in which age-τ births depend only on age-τ consumption.[35] Formally, let $f_\tau(c_\tau)$ give age-τ births as a function of age-τ consumption c_τ. Suppose that all the f_τ are strictly increasing and concave.

For any consumption vector $c = (c_1, \ldots, c_\tau)$, an indifference curve is defined by (from (6)),

$$1 \;=\; \frac{f_1(c_1)}{\Phi} + ... + \frac{f_\tau(c_\tau)}{\Phi^\tau} + \ldots + \frac{f(c_{T-1})}{\Phi^{T-1}} + \frac{f_T(c_T)}{\Phi^T}, \tag{7}$$

where Φ is constant on a particular indifference surface. A higher value of Φ corresponds to a higher indifference curve, so that consumption plan $(c'_1, \ldots, c'_T)$ is preferred to $(c_1, \ldots, c_T)$ if and only if

$$1 \;=\; \frac{f_1(c_1)}{\Phi} + \ldots + \frac{f_T(c_T)}{\Phi^T} \;<\; \frac{f_1(c'_1)}{\Phi} + \ldots + \frac{f_T(c'_T)}{\Phi^T}.$$

It follows readily that evolution's indifference surfaces over consumption bundles $(c_1, \ldots, c_\Phi)$ have the usual shape, in the sense that evolution's preferences can be described by a utility function $U(c_1, ..., c_T)$ that is strictly increasing and quasi-concave.

This gives us the beginnings of an extension from models of reproduction to models of consumption. As long as period-τ reproduction is a function only of period-τ consumption, preferences over consumption will once again be described by an exponentially-discounted sum of utilities. In practice, of course, period-τ births will depend on the entire history of consumption. At the very least, one must have consumed enough to survive until period τ in order to reproduce at that age. Period-τ births are thus implicitly a function of consumption at all preceding ages. This in turn opens the possibility that the induced preferences over consumption may exhibit complicated discounting patterns. There is much that remains to be done in terms of exploring this connection between reproduction and consumption, including especially the implications for discounting.

3.2.6 Heterogeneous Choices

We have hitherto implicitly assumed that all of our agents face the same feasible set and choose the same alternative from that feasible set. How

---

[35] See Robson, Szentes and Ianchev [128] for more involved specifications.

do we incorporate some heterogeneity into the model? In addressing this question, we keep things simple by retaining our basic framework of choice of reproductive life histories.

Suppose that each agent entering our model is randomly and independently (over time and agents) assigned one of N feasible sets, with p_n the probability of being assigned to the nth feasible set, and with $(x_1(n), \ldots, x_T(n))$ the life history chosen when faced with the nth feasible set. Some agents may find themselves in the midst of plenty and face relatively rich feasible sets, while others may face harder circumstances and more meager feasible sets. The Leslie matrix associated with this population is given by

$$\begin{bmatrix} e^{-\delta}\sum_{n=1}^{N} p(n)x_1(n) & e^{-\delta} & 0 & \ldots & 0 & 0 \\ e^{-\delta}\sum_{n=1}^{N} p(n)x_2(n) & 0 & e^{-\delta} & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ e^{-\delta}\sum_{n=1}^{N} p(n)x_{T-1}(n) & 0 & 0 & \ldots & 0 & e^{-\delta} \\ e^{-\delta}\sum_{n=1}^{N} p(n)x_{T-}(n) & 0 & 0 & \ldots & 0 & 0 \end{bmatrix}.$$

The agent's preferences can be derived from the corresponding characteristic equation, or

$$\begin{aligned} 1 &= \frac{\sum_{n=1}^{N} p(n)x_1(n)}{\Phi} + \frac{\sum_{n=1}^{N} p(n)x_2(n)}{\Phi^2} + \frac{\sum_{n=1}^{N} p(n)x_3(n)}{\Phi^3} + \ldots + \frac{\sum_{n=1}^{N} p(n)x_T(n)}{\Phi^T} \\ &= p(1)\left(\frac{x_1(1)}{\Phi} + \frac{x_2(1)}{\Phi^2} + \ldots + \frac{x_T(1)}{\Phi^T}\right) + \ldots + p(N)\left(\frac{x_1(N)}{\Phi} + \frac{x_2(N)}{\Phi^2} + \ldots + \frac{x_T(N)}{\Phi^T}\right). \end{aligned}$$

In each of these choice situations, it follows that the optimal decision is consistent with exponential discounting, where the discount rate now depends on the overall *population* growth rate. Hence, those agents facing relatively meager feasible sets will apply a discount factor seemingly higher than would be warranted from consideration of that feasible set alone, while those facing a quite rich feasible set would apply a discount factor seemingly too low. Given the discount factor, however, we would observe a collection of choices that could together be rationalized as maximizing the same exponentially discounted utility function.[36]

---

[36] One can well imagine more complicated ways in which heterogeneity might be incorporated into the model, requiring a more sophisticated model. The tools for addressing such questions are provided by the theory of structured populations, as in Charlesworth [23].

3.2.7 Nonexponential Discounting

The message to emerge from our analysis thus far is that we can expect to see agents evaluating intertemporal trades according to an exponentially-discounted utility function. Depending on one's point of view, this represents good news or bad news. On the one hand, it directs attention to the most common model of intertemporal choice in economics. At the same time, it provides little insight into departures from exponential discounting.

There are three obvious possibilities for exploring foundations of nonexponential discounting. Section 3.2.5 raises the first. Even if reproduction is discounted exponentially, the relationship between reproduction and consumption may be complicated and may induce nonexponential discounting of consumption. This possibility remains relatively unexplored.

Second, Sozou [147] and Dasgupta and Maskin [31] show that if the realization of a future consumption opportunity is subject to uncertainty, then the result can be a present bias in discounting. As illustrated by such proverbs as "a bird in the hand is worth two in the bush," the idea that one should discount uncertain prospects is quite familiar.

Sozou supposes that there is a constant hazard rate that an opportunity to consume in the future may disappear before the proposed consumption date arrives. Someone else may consume the resource beforehand, or a predator may in the meantime block access to the resource. In the absence of any additional complications, this uncertainty has a straightforward effect on the agent's behavior. Future payoffs are again exponentially discounted, with the relevant discount rate now being the sum of the death rate, population growth rate, and disappearance rate.

Sozou further assumes that the agent is uncertain about the hazard rate of consumption disappearance, updating her prior belief about this value as time passes. Suppose, for example, the agent initially compares one unit of consumption at time 0 with c units at time $t > 0$, and discounts (taking into account the likelihood that the latter will disappear before time t arrives) the latter at rate 10%. Now suppose that time $t/2$ has arrived, and the agent must again compare a unit of current (i.e., time $t/2$)) consumption with the same c units of consumption at time t. If this choice is to be meaningful, it must be the case that over the interval $\left[0, \frac{t}{2}\right]$, the future consumption opportunity did not vanish. This is good news, leading the agent to conclude that the probability of disappearance is not as high as the agent's prior distribution indicated. As a result, the agent's discount rate will now be lower than the 10% relevant at time 0.

More generally, let c_τ denote consumption at time τ. The agents in

Sozou's model apply a higher discount factor when comparing c_0 and c_1 than when comparing c_τ and $c_{\tau+1}$: if the latter choice is still relevant at time τ, then the agent will infer that the hazard rate at which consumption opportunities disappear is lower than originally suspected. As a result, the discount rate decreases as one considers choices further and further into the future, introducing a present bias into discounting.

Sozou's model will not generate preference reversals, the strikingly anomalous choices that have fueled much of the interest in present-biased preferences. In a typical preference reversal, an agent prefers $c_{\tau+1}$ from the choice $\{c_\tau, c_{\tau+1}\}$ when choosing at time 0, but then prefers c_τ when making the choice at time τ. Invoking some stationarity, the standard route to constructing a preference reversal is to assume that the agent prefers c_0 from $\{c_0, c_1\}$ at time 0 as well as prefers $c_{\tau+1}$ from the choice $\{c_\tau, c_{\tau+1}\}$; coupled with an assumption that the agent makes the choice from $\{c_\tau, c_{\tau+1}\}$ at time τ precisely as she does the choice $\{c_0, c_1\}$ at time 0. It is this latter assumption that does not hold in Sozou's model. If the choice from $\{c_\tau, c_{\tau+1}\}$ is relevant at time τ, then the agent infers that the hazard rate at which consumption opportunities disappear is not as large as originally suspected. This only reinforces the patience that prompted the agent to originally prefer $c_{\tau+1}$ from the choice $\{c_\tau, c_{\tau+1}\}$. Discount rates are thus not constant, but we would not observe the type of inconsistency in behavior that would induce the agent to take steps to restrict future choices.

In Dasgupta and Maskin [31], there is again the possibility that a consumption opportunity might disappear before it arrives, but the hazard rate at which this happens is constant and known. In the absence of any other considerations, we would then simply have constant discounting at this hazard rate (plus the relevant death and growth rates). On top of this, however, Dasgupta and Maskin add some additional uncertainty about *when* as well as whether the consumption will be realized. An opportunity to consume c_τ at time τ in fact gives the consumption at time c_τ with high probability, but with the remaining probability gives a consumption opportunity whose timing is distributed over the interval $[0, \tau]$ (all conditional on not having disappeared in the meantime). Fortuitous circumstances may bring the opportunity early.

Now consider two consumption opportunities, one promising consumption c_τ at time time τ and one promising $c_{\tau'}$ at time $\tau' > \tau$. Suppose that at time 0, the agent prefers opportunity $(c_{\tau'}, \tau')$. If this is to be the case, then we must have $c_{\tau'} > c_\tau$, since it would not be worth waiting longer for a lower reward. Now consider what happens as time passes. The dates τ and τ' at which the consumption opportunities will be realized draw nearer.

This increases the value of each option, but this effect alone does not change the relative ranking of the two consumption prospects. The probability that either one is realized is scaled upward by a common factor reflecting that an interval has passed without the consumption disappearing. The other effect is that this same interval has passed without either consumption opportunity arriving early. This decreases the value of each option, but especially decreases the value of option $(c_{\tau'}, \tau')$, since it involves the larger quantity of consumption and hence its early arrival is a relatively lucrative outcome. Thus, as time passes, the relative ranking shifts toward (c_τ, τ). If the two bundles are sufficiently closely ranked to begin with, and if the prospect of early arrival is sufficiently important, preferences will reverse to bring (c_τ, τ) into favor as time passes.

Dasgupta and Maskin's analysis thus provides us with an evolutionary account of preference reversals. At the same time, it does not give rise to the sorts of inconsistency and commitment issues that appear in behavioral models. The preference reversal as time τ draws near reflects an optimal response to the changing time-profile of the consumption opportunities. As a result, an agent would never have an incentive to preclude such reversals. Preference reversals have excited interest from behavioral economists to a large extent because people often take costly measures to avoid them. We build rigidities into our lives to ensure that currently-optimal choices are not undone by future preferences shifts. Dasgupta and Maskin's agents would welcome any preference reversals they encounter.

Dasgupta and Maskin sketch an extension of their model that gives rise to commitment issues. Very roughly speaking, they suppose that evolution has endowed people with preferences that are appropriate for the distributions of early consumption arrivals that were common over the course of our evolutionary history. Then they consider an agent facing a choice that the agent knows to involve distributions atypical of this history. An agent who simply expresses her preferences may then find herself confronted with a preference reversal which she would regard as inappropriate, given her knowledge of how the distribution of early arrivals has shifted. Given the opportunity, the agent would rationally strive to prevent such a reversal, giving rise to incentives for commitment reminiscent of behavioral models. This gives us a mismatch model of preference reversals. Must evolutionary models of preference reversals necessarily involve mismatches, or are there circumstances under which evolutionary design calls for preference reversals in the environment giving rise to that design? If the latter type of models can be constructed, is there any reason to prefer them to mismatch models? Do their implications differ? These are open and interesting questions.

The preferences emerging from the models of Sozou [147] and Dasgupta and Maskin [31] give rise to a delicate issue of interpretation. First, an essential feature of both models is that consumption opportunities are subject to uncertainty. Each model begins with the assumption that the evolutionary objective is to maximize total consumption, with discounting reflecting the uncertainty inherent in pursuing a consumption opportunity. In short, it is better to consume now rather than later because the later opportunity may disappear before it can be realized. However, the analysis of Sections 3.2.1–3.2.4 suggests that even in the absence of uncertainty (and in the absence of death), we can expect discounting, so that maximizing total consumption is not an obvious point of departure. Fortunately, building the type of considerations uncovered in Sections 3.2.1–3.2.4 into the models of Sozou or Dasgupta and Maskin appears to be straightforward.

Second, our underlying view is that evolution shapes our behavior, with preferences being an analytical tool we choose to represent this behavior. The standard approach in constructing this representation is to use preferences and feasible sets to capture different aspects of an agent's choice problem, with the feasible set describing the alternatives and constraints on the choice. In particular, the standard approach would view consumption opportunities subject to uncertainty and consumption opportunities without uncertainty as different objects, with preferences first defined in the absence of uncertainty and then extended to uncertain outcomes, perhaps via an expected utility calculation. In using discounting to capture the effects of uncertainty about consumption, the models of Sozou and Dasgupta and Maskin blur the distinction between the feasible set and preferences.

In some cases, this blurring may be precisely what is required. In particular, suppose our evolutionary model of behavior incorporates the mismatch possibility that preferences evolved in one environment but may be applied in another. If this is the case, then we must know not only the choices induced by evolution, but also the process by which these choices are induced. We thus have no alternative but to model the mechanics of the agents' decision making. It may well be that evolution has responded to some of the uncertainty in our environment by altering our discounting rather than our representation of the feasible set. Notice, however, that establishing the process by which choices are implemented is a taller order than describing the choices themselves.

An alternative possibility under which preferences may no longer exhibit exponential discounting is explored by Robson and Samuelson [125], and returns us to the distinction between idiosyncratic and aggregate risk examined in Section 3.1. We have assumed in Sections 3.2.1–3.2.6 that the uncertainty faced by the agents is idiosyncratic. It seems reasonable to imagine that aggregate uncertainty may well have been an important feature of our evolutionary environment. Periods in which the weather was harsh, food scarce, disease rampant, or predators prevalent may have a common impact on a population. What effect does this have on our analysis of time preference?

To capture the possibility of aggregate uncertainty, we assume that in each period t, a Leslie matrix $X(t)$ is drawn from a distribution over such matrices, with $X(t)$ then describing the fate of the population, in terms of both reproduction and death, during that period. A period of particularly harsh weather may be characterized by a Leslie matrix with high death rates, while a period in which food is quite plentiful may be characterized by favorable survival rates. The matrix $X(t)$ may itself contain values that are the averages of idiosyncratic uncertainty, but as before this will have no effect on the analysis.

Given an initial population $N'(0) = (N_1(0), \ldots, N_T(0))$ with $N_\tau(0)$ of agents of age τ, the population at time t is then given by (cf. (5))

$$N'(t) = N'(0)\tilde{X}(1)\tilde{X}(2)\cdots\tilde{X}(t),$$

where $\tilde{X}(t)$ is the random Leslie matrix in time t. We thus have a product of random matrices, a much less tractable object than the product of the fixed Leslie matrices arising in (5). It is not even immediately obvious that such a product has an appropriate limit. Fortunately, there are quite general theorems establishing the limiting growth rates of such products (e.g., Furstenberg and Kesten [56, Theorem 2] and Tanny [150, Theorem 7.1]), but the model is still considerably less tractable than the case of idiosyncratic uncertainty.

Aggregate uncertainty opens up all sorts of new possibilities for discounting patterns. We present here a simple example to illustrate some of these possibilities, leaving a more systematic analysis to Robson and Samuelson [125]. Suppose that there are T possible Leslie matrices, $X_1, \ldots, X_T$. Under Leslie matrix X_τ, only offspring born to parents of age τ survive, with expected offspring per parent denoted by x_τ. The Leslie matrices are drawn independently across periods and are equally likely in any given period. In each period and under every Leslie matrix, all existing agents face an idiosyncratic death risk, with death rate δ.

We thus have a rather extreme form of aggregate uncertainty, but one that significantly simplifies the resulting calculations, while driving home the point that aggregate uncertainty can lead to new results. Section 6.1 proves the following.

Proposition 1 *Almost surely,*

$$\lim_{t\to\infty} \frac{1}{t} \ln u'\tilde{X}(1)\dots\tilde{X}(t)v = \ln S + \frac{\sum_{\tau=1}^{T} \ln x_\tau}{\sum_{\tau=1}^{T} \tau}. \qquad (8)$$

Preferences are thus represented by the *undiscounted* sum of the logs of the offspring in each state. In contrast to our previous findings, there is no impatience here, no matter what the population growth rate (given by (8)) and death rate. A reduction in fertility at age τ reduces the growth rate via its effect on the term $\sum_{\tau=1}^{t} \ln x_\tau$, while the extent of this reduction does not depend upon the age in question.

We can push this example somewhat further. Suppose $T = 2$, to keep the calculations simple, and that instead of being independent across periods, the environment is drawn from a symmetric Markov process with persistence α, i.e., with probability α the environment in period t is the same as in period $t-1$, and with probability $1-\alpha$ the environment changes from period $t-1$ to period t. Section 6.1 proves:

Proposition 2 *Almost surely,*

$$\lim_{t\to\infty} \frac{1}{t} \ln u'\tilde{X}(1)\dots\tilde{X}(t)v = \tfrac{2\alpha \ln x_1 + \ln x_2}{2+2\alpha}.$$

For the case of $\alpha = 1/2$, or no persistence, we have Proposition 1's result that there is no discounting. Assuming $\alpha > 1/2$ generates impatience, while assuming $\alpha < 1/2$, so that environments are negatively correlated, generates negative discounting—the future is weighted more heavily that the present.

What lies behind the result in Proposition 1? Consider the generation of agents born at some time t, and for the purposes of this illustration only assume there is no death before age T.[37] Given the convention that only one age class reproduces in any period, these newborns all have parents of the same age, with any such age τ being equally likely, and with each parent giving rise to x_τ offspring.[38] These parents in turn all had parents of the same age, with any such age τ' being equally likely, and with each parent giving rise to x'_τ offspring. Continuing in this fashion, the number of agents born at time t is given by a product $x_\tau x_{\tau'} x_{\tau''} \dots$, where the sequence $\tau, \tau', \tau'', \dots$ identifies the age of the parents reproducing in the relevant period. Because the age to reproduce in each period is uniformly drawn from

---

[37] Since death rates are equal across ages, introducing death before age T involves only a normalization of the following calculations.

[38] It is this property that fails, vitiating the argument leading to Proposition 1, when births are not so perfectly synchronized.

the set $\{1, 2, \ldots, T\}$, over long periods of time each age will appear with very close to the same frequency in the string $\tau, \tau', \tau'', \ldots$, with that frequency being $1/T$. Hence, the number of births at time t is proportional to a power of $x_1 x_2 \ldots x_T$. In light of this, evolution will seek to maximize $\ln[x_1 x_2 \ldots x_T]$, leading to the no-discounting result. If expected offspring are equal across ages, then evolution is indifferent as to where an increment to expected offspring appears.

It is clearly an extreme assumption that only one age of parent has offspring in any given state of the environment. We present this result not for its realism, or because we would like to suggest that evolutionary models should lead us to expect that people do not discount, but to illustrate how aggregate uncertainty can lead to new and counterintuitive results. In Robson and Samuelson [125] we first show that if aggregate uncertainty bears equally on all survival rates, then we have a wedge between the rate of discounting and the sum of the growth and mortality rates. We then consider cases in which the extent of aggregate uncertainty in the environment is relatively small, unlike the model we have just presented. This reflects a belief that results emerging from models with relatively modest doses of aggregate uncertainty are a better point of departure for our analysis than models with drastic specifications of uncertainty. We present plausible, but by no means universal, conditions for aggregate uncertainty to lead to a present bias in discounting. Once again, however, this present bias leads to neither preference reversals nor a desire for commitment. The search for evolutionary foundations of preference reversals and commitment remains an important area of research.

3.2.8 Implications

Our search again turns to implications. We can start with the observation that discounting in general has nothing to do with death rates. An increase in the death rate simply induces a corresponding decrease in the growth rate (for fixed fertilities $(x_1, \ldots, x_T)$), leaving discounting unchanged. Higher fertility should thus correspond to higher discounting, holding the death rate constant, but higher death rates (holding fertility constant) should not. An attempt to verify these comparative static predictions would give rise to valuable and exciting research.

Looking a bit beyond our model, the remarks of the previous paragraph correspond to cross-population comparisons of discounting, in the sense that we would need to compare different populations whose discount factors have been adapted by evolution to their various circumstances. Suppose in contrast that we examine different types within a population. Here, the relevant terms in the discount factor are the average growth rate of the population and the death rate of the particular type in question. As a result, agents with higher death rates within a population should exhibit higher discount rates. Wilson and Daly [161] find just such a relationship.

Finally, the models suggest that evolution may more readily lead to non-exponential discounting, often in the form of a present bias, than to generate preference reversals. This suggests that experimental or empirical evidence may accordingly more readily exhibit declining discount factors than preference reversals. It is then perhaps unsurprising that some investigations do not find a great willingness to pay for the ability not to reverse preferences (e.g., Fernandez-Villaverde and Mukherji [42]).

4 Preferences over What?

Our next selection of topics takes us somewhat deeper into preferences, asking what we should expect to find as the arguments of the function u. The standard assumption throughout much of economics is that u depends only on an agent's own consumption, as in (1). At the same time, there is considerable suspicion that other factors also enter our preferences. As we have explained above, the goal is to incorporate such possibilities while still retaining some discipline in our work. This section examines three dimensions along which an evolutionary analysis is helpful.

Our guiding principle is that to understand our utility function, we must think through the constraints on what evolution can do in designing us to make good decisions. In each of the cases we describe in this section, in the absence of such constraints, we would come back to a standard utility function defined only over an individual's own consumption. However, if something prevents the construction of such a perfect utility function, then evolution may optimally compensate by building other seemingly anomalous features into our utility function. Intuitively, we have an evolutionary version of the theory of the second best.[39]

Under this approach, the analysis will be no more convincing than the

---

[39] Beginning with Lipsey and Lancaster [84], the theory of second best has become a pillar of welfare economics, noting that if some of the conditions for an optimal outcome fail, then moving closer to satisfying the remaining conditions may not improve welfare. In our context, we can first imagine a first-best or unconstrained design that would lead to evolutionary success for an agent. The idea is then that if feasibility constraints preclude implementing some features of this design, it may not be optimal to insist on all of the remaining features.

case that can made for the constraints. In this sense, Gould and Lewontin's [61] critique of evolutionary psychology recurs with some force, since one suspects that a judiciously chosen constraint will allow anything to be rationalized.

In response, before even embarking on this line of research, we should be willing to argue that it is prohibitively costly for evolution to enhance significantly our cognitive powers. Otherwise, we would expect evolution to simply have done away with whatever constraints might appear in our decision-making. Evolutionary psychologists routinely appeal to limits on our cognitive capabilities, finding evidence for these limits in the relatively large amount of energy required to maintain the human brain (Milton [93]), the high risk of maternal death in childbirth posed by infants' large heads (Leutenegger [83]), and the lengthy period of human postnatal development (Harvey, Martin and Brock [70]).

Notice that there is no question of evolution's designing us to solve some problems of inordinate complexity. The human eye and the attendant information processing is an often-cited triumph of biological engineering. Our argument requires only that evolution cannot ensure that we can solve *every* complex problem we encounter, and that she will accordingly adopt information-processing shortcuts whenever she can. "In general, evolved creatures will neither store nor process information in costly ways when they can use the structure of the environment and their operations upon it as a convenient stand-in for the information-processing operations concerned." (Clark [25, p. 64]).[40]

We should also expect to see evidence that humans often make mistakes in processing complicated information. For example, psychologists have conducted a wealth of experimental studies suggesting that people are poor Bayesians (e.g., Kahneman and Tversky [77]).

4.1 Context

This section, borrowing from Samuelson and Swinkels [134], examines one respect in which our utility seemingly depends upon more than simply what we consume, but with a perhaps somewhat unusual perspective. It is common to think of our utilities as depending not only on what we consume, but also on what we have consumed in the past, or on what others consume. Instead, we consider here the possibility that our utility also depends upon

---

[40] LeDoux [80] discusses the incentives for evolution to arm us with a mix of "hard-wired" and cognitive responses to our environment, arguing that many of our seemingly hard-wired reactions are engineered to economize on information processing.

what we could have consumed, but did *not* choose. A salad may be more attractive when the alternative is oatmeal than when it is steak, and toiling away at the office may be more bearable on a cold, cloudy day than a warm sunny day.[41]

It is no surprise, of course, that choices typically depend on the set of alternatives. Who would doubt that it is more tempting to skip work on a warm, sunny day than on a cold bitter one? There is little point in continuing if this is the extent of our insight. However, the key points of our analysis are that the presence of unchosen alternatives affects not just our choices but our preferences over those choices, and their ability to do so depends upon their salience. We may happily work in a windowless office on a brilliant spring day, but find that such work is much less satisfying when the office has a panoramic view. Knowing that one can order dessert is different than having the dessert cart at one's table. Knowing that it's nice outside is different than being able to see the sun and feel the warm breeze.[42]

As we have suggested, our evolutionary model will revolve around a constraint on evolution's ability to design agents. We assume in this case that evolution cannot equip her agents with a perfect prior understanding of the causal and statistical structure of the world. Our belief here is that the complexity of a perfect prior is simply out of reach of a trial-and-error mutation process.[43] Nor can the agents themselves be trusted to infer this information from our environment. An agent cannot learn the relationship between

---

[41] Gardner and Lowinson [57], Loewenstein [85], Mischel, Shoda and Rodriguez [95], and Siegel [142] examine the importance of salient alternatives. The possibility that preferences over objects may depend on the set from which they are chosen has attracted theoretical and experimental attention from psychologists (e.g., Tversky and Simonson [154] and Shafir, Simonson and Tversky [141]). Gul and Pesendorfer [64] present a model of such preferences centered on the assumption that resisting tempting alternatives is costly. Laibson [79] examines a model in which instantaneous utilities adjust in response to external cues. Our interest here is not so much the mechanism by which this interaction between the set of alternatives and the utility of particular alternatives is generated, but rather the question of why evolution might have endowed us with such preferences in the first place.

[42] In a similar vein, psychologists have suggested that our behavior is driven partly by a collection of utility-altering visceral urges (Loewenstein [85]). It is again straightforward to appreciate why we have urges reflecting direct evolutionary consequences such as hunger, thirst, or fatigue (Pluchik [107]). We consider here the less obvious question of why the strength of these urges can depend on the set of unchosen consequences.

[43] For example, it is difficult to randomly create an agent who knows not only that the probability of a successful birth from a random sexual encounter is about 2% (Einon, [35]), but also how this probability varies systematically with health, age, and other observable features of the mate.

specific nutrients and healthy births by trial and error quickly enough to be useful, and we certainly cannot learn quickly enough that even many generations of ample food might still be followed by famine in the next year.[44]

4.1.1 A Model

An agent in this model enters the environment and must either accept or reject an option. Accepting the option leads to a lottery whose outcome is a success with probability p and a failure with probability $1-p$. Rejecting the option leads to a success with probability q and a failure with probability $1-q$. This is the only decision the agent makes. As usual, this leaves us with a ludicrously simple evolutionary model, but one that allows us to focus clearly on the important features of the problem.

We might think of the option as an opportunity to consume and success as reproducing. The parameters p and q are random variables, reflecting the benefits of eating and the risks required to do so in any given setting. The probability of success may be either increased ($p > q$) or decreased ($p < q$) by accepting the option.

The agent is likely to have some information about the likely values of p and q. For example, the agent may know whether game is plentiful, whether food is nearby but guarded by a jealous rival, or whether a drought makes it particularly dangerous to pass up this opportunity. However, the agent is unlikely to know these probabilities precisely. We model this by assuming that the agent observes a pair of scalar signals s_p about p and s_q about q. The probabilities p and q are independent, as are the signals s_p and s_q. In addition, p and s_q are independent, as are q and s_p. Hence, each signal gives information about one (and only one) of the probabilities. We assume that s_p and s_q are informative about p and q and satisfy the monotone likelihood ratio property with respect to p and q respectively, so that (for example) $E\{p|s_p\}$ is increasing in s_p.

Evolution designs the agent to have a rule ϕ for transforming signals into estimates of the probability of success. We assume that ϕ is continuous and strictly increasing. The crucial restriction in our model—the imperfection that makes this an interesting setting for examining utility functions—is that the agent must use the *same* rule ϕ for evaluating all signals. In this simple

---

[44] This constraint is well-accepted in other areas of study. Focusing on reactions to danger, LeDoux [80, pp. 174–178] notes that evolution deliberately removes some responses from our cognitive control precisely because her prior belief is strong. "Automatic responses like freezing have the advantage of having been test-piloted through the ages; reasoned responses do not come with this kind of fine-tuning."

setting, the result is that the agent must have one process for evaluating both the signal s_p and the signal s_q, rather than a separate evaluation rule for each signal. If, for example, p and q come from different processes and with information of varying reliability, proper Bayesian updating requires that different updating rules be applied to s_p and s_q. Our assumption is that evolution cannot build this information about the prior or signal-generation process into the agent's beliefs, and hence that the agent has a single belief-formation rule ϕ.[45]

Evolution's goal is to maximize the probability of a success. In pursuit of this goal, evolution can design a utility function for the agent, with utility potentially derived both from the outcome of the agent's action and from the action itself. A success leads to an *outcome* (e.g., successful reproduction) that yields a utility of x. A failure gives the agent a utility that we can normalize to zero. In the absence of any constraints, evolution would need only these two tools. Given the agent's imperfect information process, it is potentially relevant that the *act* of accepting the option (e.g., eating the food) yields a utility of y.[46]

4.1.2 Utility

We view evolution as choosing values x and y that maximize an agent's probability of success. No generality is lost by taking $x = 1$. The question is the choice of y. If $y = 0$, then utilities are attached only to outcomes and not to actions. In this case, we would be motivated to eat not because we enjoy food, but because we understand that eating is helpful in surviving and reproducing. If y is nonzero, then actions as well as outcomes induce utility.

The optimal decision rule from an evolutionary perspective is to accept the option whenever doing so increases the probability of success, or

$$\text{accept} \quad iff \quad p - q > 0. \tag{9}$$

The agent will accept the option whenever it maximizes utility, or

$$\text{accept} \quad iff \quad y + \phi(s_p) - \phi(s_q) > 0. \tag{10}$$

---

[45] Without this restriction, the solution to the problem is again trivial. Evolution need only attach a larger utility to a success than to a failure, while designing the agent to use Bayes' rule when transforming the signals he faces into posterior probabilities, to ensure that the agent's choices maximize the probability of success.

[46] Attaching another utility to the act of rejecting the option opens no new degrees of freedom at this stage.

Consider

$$E\{p - q|\phi(s_p) - \phi(s_q) = t\}.$$

This is the expected success-probability difference $p - q$ conditional on the agent having received signals that lead him to assess this difference at t. To make our results easier to interpret, we assume throughout that the signal generating process ensures

$$\frac{dE\{p - q|\phi(s_p) - \phi(s_q) = t\}}{dt} \geq 0, \tag{11}$$

so the expected difference in success probabilities $p - q$ is weakly increasing in the agent's assessment of this difference.[47]

We then have the following characterization of the optimal utility function:

Proposition 3 *The fitness-maximizing y satisfies*

$$E\{p - q|\phi(s_p) - \phi(s_q) = -y\} = 0. \tag{12}$$

In particular, the agent's fitness is maximized by setting $y = 0$ if and only if

$$E\{p - q|\phi(s_p) - \phi(s_q) = 0\} = 0. \tag{13}$$

To see why this should be the case, we need only note that when conditions (11) and (13) hold, setting $y = 0$ ensures that the agent's choice rule (10) coincides with the (constrained) optimal choice rule (9). There is then no way to improve on the agent's choices and hence setting $y = 0$ is optimal. More generally, let us fix a value of y and then consider the expectation $E\{p - q|\phi(s_p) - \phi(s_q) = -y\}$, which is the expected difference in success probabilities at which the agent is just indifferent between accepting and rejecting the option. If this expectation is positive, then the expected probability of success can be increased by increasing y, and if this expectation is negative, then the expected probability of success can be increased by decreasing y, giving the result.

From (13), if the agent interprets his signals correctly, then there is no evolutionary value in attaching utilities to actions. The agent will make appropriate choices motivated by the utility of the consequences of his actions.

---

[47] This is an intuitive assumption and it is easy to find either examples in which it is satisfied or sufficient conditions for it to hold, but it is *not* simply an implication of our monotone-likelihood-ratio-property assumption.

The agent will still sometimes make mistakes, but without better information there is no way to eliminate these mistakes or improve on the expected outcome.

From (12), if the agent does not interpret his signals correctly, then evolution will attach utilities to his actions in order to correct his inferences at the *marginal* signal, i.e., at the signal at which the expected success probabilities are equal. The agent must be indifferent $(y+\phi(s_p)-\phi(s_q)=0)$ when his signal would lead a perfect Bayesian to be indifferent $(E\{p-q|\phi(s_p)-\phi(s_q)=-y\}=0)$.

An initial expectation might be that evolution should attach utilities only to the things evolution "cares" about, or outcomes, rather than actions. As Proposition 3 confirms, we have rendered this suboptimal by giving the agent an unreliable understanding of how actions translate into outcomes. Evolution then compensates by attaching utilities to actions. One might then expect utilities to reflect the *average* evolutionary value of the various actions. Those that often lead to success should get large utilities, those that are less productive should have smaller utilities. However, Proposition 3 indicates that this intuition need not hold, for two reasons. First, we can expect utilities to be attached to actions only to the extent that agents sometimes misunderstand the likelihoods of the attendant outcomes. If the outcomes are correctly assessed, then actions, no matter how valuable, need receive no utility. Optimal utilities thus reflect not the evolutionary value of an action, but the error the agent makes in assessing that evolutionary value. Second, one might think that fitness would be maximized by a utility function that corrected this error *on average.* As (12) makes clear, what counts is the error the agent makes in the marginal cases where he is indifferent between two actions.

We illustrate by constructing an example in which the agent on average overestimates the value of accepting the option, but evolutionary fitness is nonetheless improved by setting $y>0$, pushing him to accept the option more than he otherwise would. Let

$$E\{p-q|\phi(s_p)-\phi(s_q)=t\}=a+bt,$$

with $a>0$ and $b>0$. Solving (12), the optimal utility is

$$y=\frac{a}{b}. \tag{14}$$

Assume that $\phi(s_p)-\phi(s_q)$ is large on average and that $b<1$. Because $\phi(s_p)-\phi(s_q)$ is on average large and $b<1$, the agent on average overestimates the value of the option. However, since $y = \frac{a}{b} > 0$, the agent's fitness is maximized by pushing the agent even more toward acceptance. We see here the importance of the agent's marginal beliefs: When $\phi(s_p) - \phi(s_q) = -\frac{a}{b}$ (so that $E\{p - q | \phi(s_p) - \phi(s_q)\} = 0$), the agent *underestimates* the relative value of the option (thinking it to be negative), even though he overestimates it on average.

It follows from (14) that, as one might expect, a choice with a large expected value (large a) will tend to have a large utility. It is thus no surprise that we have a powerful urge to flee dangerous animals or eat certain foods. However, there is also a second effect. The smaller is b, the larger is y. The point is that the less informative is the agent's information, holding fixed his average assessment, the more negative is the relevant marginal signal. When b is near zero, evolution effectively insists on the preferred action. While blinking is partly under conscious control, our utility functions do not allow us to go without blinking for more than a few seconds. It would seem that we are unlikely to have reliable information suggesting that this is a good idea.

4.1.3 Choice-Set Dependence

We have reached a point where evolution might optimally attach utilities to actions, but have said nothing about how utilities might depend upon the set of salient alternatives. In this section, we show how a setting where the agent makes different mistakes in different contexts creates evolutionary value for a utility function that depends on things that have *no* direct impact on evolutionary success. Rather, their role is to tailor utility more closely to the specific informational context at hand. How any given feature optimally affects utility thus depends both on its direct evolutionary impact *and* how it correlates with errors in information processing.

Suppose that the environment may place the agent in one of two situations. The success probability when rejecting the option is q in either case, with success probability p_1 and p_2 when accepting the option in situations 1 and 2. The corresponding signals are s_q, s_{p_1} and s_{p_2}. We initially assume that, as before, the agent derives a utility of 1 from a success, 0 from a failure, and utility y, *the same value in both situations*, from the act of accepting the option.

For example, suppose that in situation 2, accepting the option entails an opportunity to eat a steak. As we have shown, evolution optimally attaches a utility y to steak satisfying

$$E(p_2 - q | \phi(s_{p_2}) - \phi(s_q) = -y) = 0.$$

Now suppose that in situation 1, accepting the option entails eating a steak at the end of a hunting trip. The agent is likely to have quite different sources of information about these two situations and thus to make quite different errors in processing this information. In particular, the hunter may have an idea of what hazards he will face on the hunting trip before achieving consumption and how these will affect the probability p_1. Only coincidentally will it then be the case that $E(p - q|\phi(s_p) - \phi(s_q) = -y$, steak on hand) equals $E(p - q|\phi(s_p) - \phi(s_q) = -y$, steak to be hunted). But if these two are not equal, the agent's expected fitness can be increased by attaching different utilities to accepting the option in the two situations.

How can evolution accomplish this? One possibility is to attach utilities to more actions. The agent can be given a taste for meat, a disutility for the physical exertion of hunting, and a fear of the predators he might encounter. However, there are limits to evolution's ability to differentiate actions and attach different utilities to them—what it means to procure food may change too quickly for evolution to keep pace—and the set of things from which we derive utility is small compared to the richness of the settings we face. As a result, evolution inevitably faces cases in which the same utility is relevant to effectively different actions. This is captured in our simple model with the extreme assumption that y must be the same in the two situations. The critical insight is then that the agent's overall probability of success can be boosted if utility can be conditioned on some other reliable information that is correlated with differences in the actions.

Assume that in situation 2, a utility of z can be attached to the act of *foregoing* the option. We say that an option with this property is *salient.* In practice, an option is salient if its presence stimulates our senses sufficiently reliably that evolution can tie a utility to this stimulus, independently of our signal-processing.[48] In our example, the presence of the steak makes it salient in situation 2. The question now concerns the value of z. If fitness is maximized by setting $z \neq 0$, then there is evolutionary advantage to tailoring the utility gradient between accepting and rejecting the option to the two situations, and we have "choice-set dependence." Only if $z = 0$ do we have a classical utility function.

Proposition 4 *The optimal utility function* (x, y, z) *does **not** exhibit choice-*

---

[48] The importance of salient alternatives is well studied by psychologists (Gardner and Lowinson [57], Mischel, Shoda and Rodriguez [95], Siegel [142]) and is familiar more generally—why else does the cookie store take pains to waft the aroma of freshly-baked cookies throughout the mall?

set dependence (sets $z = 0$) if and only if there exists t^* such that

$$E\{p_1 - q|\phi(s_{p_1}) - \phi(s_q) = t^*\} = E\{p_2 - q|\phi(s_{p_2}) - \phi(s_q) = t^*\} = 0. \quad (15)$$

To see why this is the case, we note that if (15) holds, then the agent's estimates of the success probabilities in the two situations he faces are equally informative at the relevant margin. Setting $z = 0$ and $y = -t^*$ then ensures that (12) holds in both situations, and there is thus no gain from choice-set dependence. Conversely, suppose that the agent's beliefs are differentially informative in the two situations (i.e., (15) fails). Then fitness can be enhanced by attaching different utility subsidies in the two situations. This can be accomplished by choosing y to induce optimal decisions in situation 1 and $y - z$ (and hence $z \neq 0$) to induce optimal decisions in situation 2. The result is choice-set dependence.

For example, using choice-set dependence to boost the relative attractiveness of steak when it is available ($z < 0$), in contrast to simply increasing the utility of steak across the board (increasing y), might reflect a situation in which evolution finds it beneficial to grant substantial influence to the agent's beliefs about the consequences of production, while allowing less influence to his beliefs about consumption.

4.1.4 Implications

Our model of the evolution of choice in the face of coarse priors tells us that evolution will generally find it useful to exploit choice set dependence. Anyone who has ever said, "Let's put these munchies away before we spoil our dinner," or more generally "I don't keep junk food in the house because I know I'd eat too much if I did," has practical experience with choice-set dependence. Best of all is to be without the temptation of a pantry full of sinfully delicious snacks. Once they are there, eating is the preferred choice. Worst of all is looking at the food, constantly knowing it is there, without indulging.[49] In essence, such an individual is engaged in the sort of evolutionary conflict described in Section 2.3. If the agent's utility function perfectly captured the evolutionary goals it was designed to pursue, there would be no conflict, but the same complexity that forces evolution to resort to the device of a utility function also makes it difficult to design a perfect utility function. As a result, the utility function sometimes pulls

---

[49] Thaler [151, p. xv] tells of a request to put tempting munchies aside, coming from a group of people seemingly well acquainted with decision theory, and explains it with much the same preferences.

the individual in a direction unintended by evolution. This gives rise to a potentially intricate game, in which evolution resorts to devices such as context dependence to reinforce her desired ends, while the agent seeks refuge in devices such as hiding (or not buying) the junk food.

Which alternatives are salient in any given context is again the result of evolution. As it turns out, a sizzling steak is salient while a steak in the grocer's freezer is not. Potato chips on the table are salient; those in the pantry are less so. What is salient reflects both the technological constraints faced by evolution and the incremental value of tailoring utility to specific contexts.

Choice-set dependence can give rise to internal conflict and problems of self control. For example, suppose the agent begins by choosing between an unhealthy but gratifying meal and a diet meal. Situation 1 corresponds to a lonely meal at home, with a refrigerator full of health food and nary an ounce of fat in sight. Situation 2 corresponds to a steakhouse with a supplementary dieter's menu. Suppose that evolution has designed our preferences so that the act of choosing steak is subsidized when it is salient. Then the agent may prefer situation 1 even if there is some cost in choosing situation 1, in order to ensure that he rejects the steak.

Economists have recently devoted considerable attention to issues of self control, with present-biased preferences being a common route to self-control problems. Our best intentions to reap the benefits of a healthy diet may come to nothing if our preferences continually put excessive weight on the immediate gratification of the dessert tray. It is accordingly interesting to note that choice-set dependence has implications for self control beyond those of present bias. First, difficulties with self control can arise without intertemporal choice. One can strictly prefer junk food that is hidden to that which is exposed, knowing that one will find it painful to resist the latter, all within a span of time too short for nonstandard discounting to lie behind the results. More importantly, because our utility for one choice can be reduced by the salient presence of another, it may be valuable to preclude temptations that one *knows* one will resist. Someone who is certain she will stick to a diet may still go to some lengths not to be tempted by rich food.

When gut instincts and dispassionate deliberations disagree, the "rational" prescription is to follow one's head rather than one's heart. In our model, a strong utility push in favor of an action indicates either that the action has been a very good idea in our evolutionary past or that this is a setting in which our information has typically been unreliable. There is thus information in these preferences. The truly rational response is to ask how much weight to place on the advice they give.

4.2 Status

We now return to the consideration of status, on which we touched briefly in Section 3.1.2. The concept of status runs throughout our ordinary lives. We readily categorize people as being of high status or low status, and talk about actions as enhancing or eroding status.

We will examine a particular, narrow view of status as arising out of relatively high consumption. People's preferences often appear to depend not only on their own consumption, but also on the consumption of others, so much so that "keeping up with the Joneses" is a familiar phrase. Frank [47], Frey and Stutzer [50, 51], and Neumark and Postlewaite [98] highlight the importance of such effects, while the suggestion of a link between desired consumption and one's past consumption or the consumption of others is an old one, going back to Veblen [156] and Duesenberry [34].

There are two basic approaches to explaining such relative consumption effects. One retains the classical specification of preferences, building a model on the presumption that people care directly only about their own consumption. However, it is posited that some resources in the economy are allocated not via prices and markets but according to status. In addition, it is supposed that one attains status by consuming more than do others, perhaps because the ability to do so is correlated with other characteristics that are important for status. A flashy sports car may then be valued not only for its acceleration, but also for its vivid demonstration that the driver has spent a great deal of money. Tuna may taste better than caviar, but fails to send the same signal. The resulting behavior will be readily rationalized by preferences in which people care about their consumption and about how their consumption relates to that of others. For example, Cole, Mailath and Postlewaite [26] construct a model in which competition for mates induces a concern for status, around which a subsequent literature has grown.

The second alternative explanation is that evolution has directly embedded a concern for status into our preferences. We focus on this second possibility here, both because it is relatively unexplored and because it naturally suggests links to evolutionary foundations. As usual, our models of this possibility evolve around some constraint on evolution's ability to shape behavior. We consider two possible sources of relative consumption effects, arising out of two such constraints.

4.2.1 Information and Relative Consumption

Our first examination of relative consumption effects emphasizes information considerations, and ultimately hinges on an imperfection in information processing. The basic idea here is that relative consumption effects may have been built into our preferences as a means of extracting information from the behavior of others. We present a simple model of this possibility here, expanded and examined more thoroughly in Samuelson [133] and Nöldeke and Samuelson [99].

The idea that one can extract information from the actions of others is familiar, as in the herding models of Banerjee [7] and Bikhchandani, Hirshleifer and Welch [14]. In our case, agents observe their predecessors through the filter of natural selection, biasing the mix of observations in favor of those who have chosen strategies well-suited to their environment. An agent's observed behavior thus mixes clues about the agent's information with clues about his evolutionary experience, both of which enter the observer's inference problem. The problem then resembles that of Banerjee and Fudenberg [8] and Ellison and Fudenberg [36, 37] more than pure herding models.

At the beginning of each period $t = 0, 1, \ldots$, the environment is characterized by a variable $\theta_t \in \{\underline{\theta}, \overline{\theta}\}$. The events within a period proceed as follows:

1. Each member of a continuum of surviving agents gives birth, to the same, exogenously fixed, number of offspring. Each offspring is characterized by a parameter ϵ, with the realized values of ϵ being uniformly distributed on $[0, 1]$.

2. Each newborn observes n randomly selected surviving agents from the previous generation, discerning whether each chose action $\underline{z}$ or $\overline{z}$.

3. All parents then die. Each member of the new generation chooses an action $z \in \{\underline{z}, \overline{z}\}$.

4. Nature then conducts survival lotteries, where $h : \{\underline{z}, \overline{z}\} \times [0, 1] \times \{\underline{\theta}, \overline{\theta}\} \to [0, 1]$ gives the probability that an agent with strategy z and characteristic ϵ survives when the state of the environment is θ. Again, we assume no aggregate uncertainty.

5. Nature draws a value $\theta_{t+1} \in \{\underline{\theta}, \overline{\theta}\}$.

We interpret the actions $\underline{z}$ and $\overline{z}$ as denoting low-consumption and high-consumption lifestyles. The survival implications of these actions depend upon individual characteristics and the state of the environment. Some agents may be better-endowed with the skills that reduce the risk of procuring consumption than others. Some environments may feature more plentiful and less risky consumption opportunities than others. These effects appear in the specification of the survival probabilities $h(z,\epsilon,\theta)$, given by

$$\begin{aligned} h(\underline{z},\epsilon,\theta) &= \frac{1}{2} \\ h(\overline{z},\epsilon,\overline{\theta}) &= \frac{1}{2}+b(\epsilon-q) \qquad (16) \\ h(\overline{z},\epsilon,\underline{\theta}) &= \frac{1}{2}+b(\epsilon-(1-q)), \qquad (17) \end{aligned}$$

where $0<q<1/2$ and, to ensure well-defined probabilities, $0<b<1/(2(1-q))$. The low-consumption action $\underline{z}$ yields a survival probability of $\frac{1}{2}$, regardless of the agent's characteristic or state of the environment. The high-consumption action $\overline{z}$ yields a higher survival probability for agents with higher values of ϵ and yields a higher survival probability when the state is $\overline{\theta}$.

The environmental parameter θ follows a Markov process, retaining its current identity with probability $1-\tau$ and switching to its opposite with probability $\tau<\frac{1}{2}$.

An agent's strategy identifies an action as a function of the agent's characteristic ϵ and information. Strategies (but not characteristics or actions) are heritable and are thus shaped by natural selection.

Our interest concerns cases in which fluctuations in the state θ are not perfectly observed by the agents and are sufficiently transitory that Nature cannot observe them.[50] It follows from the monotonicity of (16)–(17) that an optimal strategy must take the form of a cutoff $\epsilon^*(\cdot)$, conditioned on the agent's information, such that action $\overline{z}$ is chosen if and only if $\epsilon>\epsilon^*(\cdot)$.

Let ψ_t be the proportion of strategy $\overline{z}$ among those agents who survived period $t-1$. Then a period-t new-born observes $\overline{z}$ on each survivor draw with probability ψ_t and observes $\underline{z}$ with probability $1-\psi_t$. Let $\Psi_{\mathcal{E}}(\psi_t,\theta_t)$ be the proportion of surviving period-t agents who chose $\overline{z}$, given that (*i*) these agents, as new-borns, drew observations from the distribution described by ψ_t, (*ii*) the period-t state of the environment relevant for Nature's survival lotteries is θ_t, and (*iii*) every agent's decision rule is given by the decision

---

[50] If the state θ can be observed, then evolution faces no constraints in designing strategies to maximize the survival probabilities given by (16)–(17), and observations of the previous generation are irrelevant for behavior.

$\mathcal{E} = \{\epsilon^*(n), \ldots, \epsilon^*(0)\}$. We can describe our system as a Markov process (ψ_t, θ_t) defined on the state space $[0,1] \times \{\underline{\theta}, \overline{\theta}\}$. Letting Θ denote the transition rule governing the state θ, $(\Psi_{\mathcal{E}}, \Theta)$ denotes the transition rule for the process (ψ_t, θ_t), where:

$$\begin{aligned} \psi_{t+1} &= \Psi_{\mathcal{E}}(\psi_t, \theta_t) \\ \theta_{t+1} &= \Theta(\theta_t). \end{aligned}$$

The optimal strategy $\epsilon^*(\cdot)$ maximizes

$$\int_{\Theta \times \Psi} \rho(\theta, \psi) \ln \left(\int_K f(k|\theta, \psi) p(\epsilon^*(k), \theta) dk \right) d\theta d\psi, \tag{18}$$

where ρ is the stationary distribution over states $(\theta, \psi) \in [0,1] \times \{\underline{\theta}, \overline{\theta}\}$, f is the distribution over the number (k) of $\overline{z}$ agents observed when sampling the previous generation (given the state (θ, ψ)), and p is the probability that an agent characterized by decision rule ϵ^* (i.e., chooses $\overline{z}$ if and only if $\epsilon > \epsilon^*$) survives in state θ. Notice in particular the ln that appears in this expression. The fluctuating state of the environment subjects the agents to aggregate uncertainty. This objective is then the adaption of (3) to this somewhat more complicated setting.

The key question in characterizing an optimal strategy is now the following: if the agent observes a relatively large value of k, is the environment more likely to be characterized by $\underline{\theta}$ or $\overline{\theta}$? Let $\rho(\overline{\theta}|k)$ be the posterior probability of state $\overline{\theta}$ given that an agent has observed k agents from the previous generation choosing $\overline{z}$. These updating rules are an equilibrium phenomenon. The expectation is that an agent observing more instances of high consumption will think it more likely that the state is $\overline{\theta}$ and hence be more willing to choose high consumption, i. e., that $\epsilon^*(k)$ should be decreasing in k. We say that a strategy $\{\epsilon^*(n), \ldots, \epsilon^*(0)\}$ is *admissible* if it exhibits this property.

Let the function $\rho_{\mathcal{E}}(\overline{\theta}_t|k, t)$ give the probability that the state in time t is $\overline{\theta}$, given a time-t observation of k values of $\overline{\theta}$. The role of k in this probability balances two considerations—the extent to which an observation of a large k indicates that the previous-period state was relatively favorable for strategy $\overline{z}$ (i.e., was $\overline{\theta}$), and the probability that the state may have changed since the previous period. Samuelson [133] proves:

Lemma 5 *There exists a value* $q^* \in (0, \frac{1}{2})$ *such that for any* $q \in (q^*, \frac{1}{2})$ *and any admissible* $\mathcal{E}$, *there exist probabilities* $\rho_{\mathcal{E}}(\overline{\theta}|k)$ $(k = 0, \ldots, n)$ *satisfying,* *for all initial conditions,*

$$\lim_{t\to\infty} \rho_{\mathcal{E}}(\overline{\theta}_t|k,t) = \rho_{\mathcal{E}}(\overline{\theta}|k).$$

The $\rho_{\mathcal{E}}(\overline{\theta}|k)$ *satisfy* $\rho_{\mathcal{E}}(\overline{\theta}|k+1) > \rho_{\mathcal{E}}(\overline{\theta}|k)$.

The restriction that $q > q^*$ ensures that the population can never get too heavily concentrated on a single action, either $\overline{z}$ or $\underline{z}$. This in turn ensures that changes in the environmental state are reflected relatively quickly in the observed distribution of actions, and hence that the latter is informative.[51]

The inequality $\rho_{\mathcal{E}}(\overline{\theta}|k+1) > \rho_{\mathcal{E}}(\overline{\theta}|k)$ indicates that observations of high consumption enhance the posterior probability that the state of the environment is $\overline{\theta}$. This is the foundation of relative consumption effects.

An equilibrium is a specification of $\mathcal{E}$ that is optimal in the induced stationary state. Hence, in defining an equilibrium, we use the limiting probabilities $\rho_{\mathcal{E}}(\overline{\theta}|k)$ to evaluate the payoff of a strategy. This reflects an assumption that the process governing the state of the environment persists for a sufficiently long time that (*i*) evolution can adapt her agents to this process, and (*ii*) the limiting probabilities $\rho_{\mathcal{E}}(\overline{\theta}|k)$ are useful approximations for evolution of the information-updating problem facing the agents. Nöldeke and Samuelson [99] show that:

Proposition 6 *There exists* $q^* \in (0, \frac{1}{2})$ *and* $\tau^* > 0$ *such that for any* $q \in (q^*, \frac{1}{2})$ *and* $\tau \in (0, \tau^*)$, *an equilibrium with an admissible strategy* $\{\epsilon^*(n), \ldots, \epsilon^*(0\}$ *exists. In any such equilibrium,* $\epsilon^*(k+1) < \epsilon^*(k)$.

Agents are more likely to choose high consumption, i.e., choose $\overline{z}$ for a wider range of ϵ, when k is large. Observations of high consumption, by increasing the expectation that the environment is in a state favorable to high consumption, increase an agent's propensity to choose high consumption. A revealed preference analysis of behavior would thus uncover relative consumption effects, in which agents optimally exploit information by conditioning their consumption on observations of others' consumption.

It is important to note that an agent's survival in this model depends only on the agent's own consumption. The route to genetic success is to choose optimal consumption levels, regardless of others' choices. The consumption

---

[51] To see how this could fail, consider the extreme case of $q = 0$. In this case, it is possible that virtually the entire population chooses $\underline{z}$. A change from state $\underline{\theta}$ to $\overline{\theta}$ will then not produce a noticeable change in the distribution of actions for an extraordinarily long time, causing this distribution to be relatively uninformative.

levels of others are relevant only because they serve as valuable indicators of environmental information that neither the agents nor Nature can observe.

There are many ways Nature could induce the optimal behavior characterized by Proposition 6, from hard-wired stimulus-response machines to calculating agents who understand Bayes' rule and their environment and who make their decisions so as to maximize the expected value of a utility function defined in terms of only their own consumption. Our argument thus far accordingly provides no reason to believe that relative consumption effects are built directly into preferences, and no reason why we should care about which of the many observationally-equivalent methods Nature might have chosen to implement such behavior.

The next step in the argument returns us to the observation that Nature faces a variety of obstacles in inducing behavior that will maximize expected utility. Suppose that in addition to the number k of preceding agents observing high consumption, the agent also observes a signal ξ that is more likely to take on high values when the environment is $\overline{\theta}$. Suppose also that the agent does not process this signal perfectly. In Samuelson [133], this imperfect-information processing assumption is made operational by assuming that the agent observes an informative signal ξ, as well as an uninformative signal ζ, but does not recognize this distinction, instead simply processing all signals as if they were informative. Recognizing that both ξ and ζ play a role in the agent's information, evolution finds the agent's information less informative than does the agent. She thus reduces the sensitivity of the agent's actions to his information. This reduced sensitivity can be accomplished by a utility function that discourages the agent from straying too far from a target action $\hat{\epsilon}(k)$ that depends upon the agent's observation of others' consumption. In particular, evolution can make the agent's utility depend upon his value of ϵ, his action ($\underline{z}$ or $\overline{z}$), and the number k of high-consumption agents observed in the previous period (the relative consumption effect). Consider a value ϵ^* and the posterior belief $\hat{\rho}_{\mathcal{E}}(\theta|k, \xi, \zeta)$ that would make the cutoff ϵ^* optimal given perfect information processing. Given that the agent is sometimes responding to an uninformative signal, evolution now has an incentive to boost the agent's marginal utility at ϵ^* above zero (i.e., $\epsilon^* < \epsilon(k)$) if the agent has received a large signal convincing him that $\overline{\theta}$ is quite likely; or depressed below zero (i.e., $\epsilon^* > \hat{\epsilon}(k)$), if the agent has received a small signal. Evolution thus requires that the agent observe more persuasive information than would be the case with errorless information processing before straying too far from a consumption strategy that makes high consumption more likely when more instances of high consumption have been observed. Evolution accomplishes this by not only inducing the agent's behavior to respond to the behavior of others, but by using the ability to make the agent's utility respond to the behavior of others.

We now have relative consumption effects built directly into preferences, in order to induce relative consumption effects in behavior. Notice that the case for the preference effect is somewhat more tenuous than for the behavioral effect. We can expect relative consumption effects in behavior whenever agents face environmental uncertainty. Relative consumption effects in preferences are one solution to a particular constraint in Nature's design problem. However, the general principle remains that if Nature cannot ensure the agent processes information perfectly, then she will find it advantageous to compensate by manipulating other features of the agent's decision-making apparatus, with relative consumption effects in preferences being one possible result.

4.2.2 Adaptive Utility and Relative Consumption

Our next approach views relative consumption effects as arising out of constraints on the technology for translating choices into utilities that evolution can build into her agents. This line of work, beginning with Robson [121, pp. 17–19], brings us back to an old question in economics—is utility reasonably viewed as a cardinal or ordinal concept?

The concept of cardinal utility traces back to the English philosopher and lawyer Jeremy Bentham [12]. Bentham believed that utility derived from pleasure or pain, and proposed to make judgments about policy by summing these utilities across the individuals involved. The result was his maxim "the greatest good for the greatest number," which, as Paul Samuelson is said to have remarked, has too many "greatests" in it to be implementable. Whatever the value of the maxim, the point of view was clear, namely that utility was a physical process whose properties we could discover and whose nature would provide clues as to how and why people make choices.

The view that utility is a cardinal notion, perhaps based on some measurable concept of pleasure, raises a number of awkward questions. Perhaps as a result, subsequent economists pared back the notion of utility to take refuge in an ordinal interpretation. In the context of consumer theory, it was realized that utility simply did not need to be cardinal—one needed only indifference curves and an appropriate set of labels. That such stripping down was philosophically a good idea was justified by an appeal to "Occam's Razor." Although matters are less cut-and-dried in the original context of welfare theory, most economists also became skeptical of interpersonal comparisons based on cardinal utility, often settling finally for a weak welfare criterion that is independent of any such comparisons—Pareto efficiency. This is associated with a clear minimal view of utility, as simply a description of choice, devoid of any physical or extraneous causal features.

This reliance on ordinal utility, while convenient from both a conceptual and technical point of view, has begun to falter in response to recent work in psychology and behavioral economics. As this work has illustrated an ever-more complicated and subtle array of choice behavior, it has been natural to seek explanations in the process by which these choices are made, in the course of which utility once again often plays the role of a mechanism rather than description.[52] For example, psychologists discuss how a burst of intense pleasure stems from a positive outcome, such as winning the lottery, but this pleasure subsides fairly quickly, with the winner ending up feeling only slightly better than before winning. Analogously, the intense sadness that arises from a negative outcome, such as becoming the victim of a crippling accident, tends to fade away, so that one ends up feeling only somewhat worse than before the accident.[53] In both cases, the dominant effect is that if you were happy before, you will be happy now; if you were miserable before, you will be miserable now. Taken at face value, these findings seem to suggest that people should not particularly mind running the risk of a catastrophic accident and should not buy lottery tickets. Why take precautions to avoid only a slight loss, or incur costs in search of a slight gain? But people do try to avoid being maimed and do buy lottery tickets.

Putting these considerations together, we consider here a model of utility with three features. Utility is a physical process that translates actions and choices into rewards, typically described as pleasure. In addition, these rewards are adaptive. Whether an experience makes you happy or sad depends on what you were expecting, on what you had before, on what those around you are receiving. Moreover, this adaption is not always perfectly anticipated. We buy lottery tickets because we anticipate the resulting utility boost, without recognizing that it will be adapted way, and avoid accidents for similar reasons.

---

[52] Recent experiments have provided fascinating evidence of the link between utility and chemical processes in the brain. See, for example, Zaghloul, Blanco, Weidemann, McGill, Jaggi, Baltuch and Kahana [168].

[53] Attention was drawn to this phenomenon by Brickman, Coates and Janoff-Bulman's [18] study of lottery winners and paraplegics, and has become the subject of a large literature. See Loewenstein and Schkade [87] and Frederick and Loewenstein [48] for introductions and Gilbert [59] for a popular account.

It will be helpful to begin with an analogy. Consider an old-fashioned, analog voltmeter, with a needle that is moved along a scale by an electrical current. To get an accurate reading from a voltmeter, one must first estimate the range into which the unknown voltage falls. If the range is set too high and the resulting voltage is in fact quite low, the needle hardly budges and the voltmeter produces no useful information. If the range is set too low, the meter self-destructs as the needle pegs against its upper end and the unexpected surge of current burns out the meter. Only if the range is set right can you obtain useful information. The problem is that the voltmeter, like all real measuring devices, has limited sensitivity.

The suggestion here is that one might think similarly about utility. The ultimate rewards that motivate our choices are provided by chemical flows in our brain. There are limits to the strength of these flows. In addition, we are likely to have limited perceptual discrimination, being unable to reliably tell the difference between roughly similar perceptual stimuli.

Consider the following example. An individual must choose between two lotteries over real numbers, with larger outcomes being better than smaller ones. Each lottery is an independent draw from the same known continuous cumulative distribution function F. The individual must choose a lottery *after* the draws are made. The choice then seems stunningly simple—there is no need to worry about expected values, or risk, or anything else. Just pick the larger number. However, suppose that the individual can only perceive whether each realization is above or below some threshold c. Evolution creates incentives to make the right choice by attaching hedonic utilities to the perceived outcomes, being high when an outcome above c is selected and otherwise low. If the outcomes of both lotteries lie above or both lie below c, choice is made randomly, so that with probability 1/2 the individual makes a mistaken choice, failing to choose the larger value.

What value of c minimizes the probability of error, given the distribution F from which choices are made? This probability of error is

$$\begin{aligned} PE(1) &= (1/2)\Pr\{x_1, x_2 < c\} + (1/2)\Pr\{x_1, x_2 > c\} \\ &= (1/2)(F(c))^2 + (1/2)(1 - F(c))^2 \\ &= (1/2)y^2 + (1/2)(1 - y)^2, \end{aligned}$$

where x_1 and x_2 are the outcomes of the two lotteries and $y = F(c)$. This is a convex function. The first-order condition for this minimization problem is

$$\frac{dPE(1)}{dy} = y - (1 - y) = 0,$$

so that one should choose c so that $y = F(c) = \frac{1}{2}$. Hence, it is optimal to choose c to be the median of the distribution described by F. In particular, it is optimal to set a threshold that adapts to the circumstances in which it is to be used, as captured by F.

We view this simple example as a metaphor for the problem evolution faces when designing utility functions. In the absence of any constraints, evolution would simply give the agent the utility function x, and would be confident of optimal decisions. An ordinal view of utility would be perfectly adequate. The view of utility as arising out of a process for transforming choices into rewards introduces constraints, in that values of x that are quite similar might induce sufficiently similar rewards that the agent sometimes ranks them incorrectly.[54] We have taken this to the extreme here of assuming that the agent can only distinguish high from low. This in turn gives rise to a design problem. If the utility function is going to give rise to imperfections, then evolution will want to influence and allow for those imperfections. This gives us our first look at the first of the three features we would like to build into our model of adaptive utility.

Before looking for the next feature, namely the adaptive part, we pause to elaborate on our first example. There is clearly a long way to go from this example to models of utility functions. To begin, the probability of error is not the most convincing objective here. After all, some errors involve a very large gap between the x that is chosen and the optimal x, and some involve a very small gap. A more plausible objective would be to identify fitness with x and then maximize the expected value of the x that is received.[55] Now the value of the threshold c should be set at the mean of the distribution rather than the median. Having done this, an obvious next question is to ask what happens if the agent is somewhat more sophisticated than being able to identify only a single threshold for the value of x.

Netzer [97] examines this problem further, considering the case in which the individual maximizes the expected payoff and has an arbitrary number of perception thresholds available. We will continue here with the illustrative and more tractable problem of minimizing the probability of error, now considering the more general case in which the individual has N threshold

---

[54] The psychology literature is filled with studies documenting the inability of our senses to reliably distinguish between small differences. For a basic textbook treatment, see Foley and Matlin [44].

[55] The identification of fitness with x is relatively innocuous, in the sense that, if fitness were a monotonically increasing function of x, we could easily find the cumulative distribution function over fitness that is implied by the given distribution over x. This does not make a significant qualitative difference.

values

$$c_1 < c_2 < ... < c_N.$$

The probability of error is now

$$\begin{aligned} PE(N) &= (1/2)(F(c_1))^2 + ... + (1/2)(F(c_{n+1}) - F(c_n))^2 + ... + (1/2)(1 - F(c_N))^2 \\ &= (1/2)(y_1)^2 + ... + (1/2)(y_{n+1} - y_n)^2 + ... + (1/2)(1 - y_N)^2, \end{aligned}$$

where $y_n = F(c_n)$ for $n = 1, ..., N$. This is again a convex function of $(y_1, ..., y_N)$ so that satisfying the first-order conditions is still necessary and sufficient for a global minimum. These first-order conditions are

$$\begin{aligned} \frac{\partial PE(N)}{\partial y_1} &= 0 \text{ so } y_2 - y_1 = y_1 - 0 \\ \frac{\partial PE(N)}{\partial y_n} &= 0 \text{ so } y_{n+1} - y_n = y_n - y_{n-1}, \text{ for } n = 2, ..., N-1 \\ \frac{\partial PE(N)}{\partial y_N} &= 0 \text{ so } 1 - y_N = y_N - y_{N-1}. \end{aligned}$$

Hence, the solution is

$$y_1 - 0 = k, \; y_{n+1} - y_n = k, \text{ for } n = 2, ..., N-1 \text{ and } 1 - y_N = k.$$

It must then be that $k = 1/(N+1)$, so that

$$y_n = F(y_n) = n/(N+1), \text{ for } n = 1, ..., N.$$

For example, if $N = 9$, the thresholds should be at the deciles of the distribution.

What is the probability of error $PE(N)$ when the thresholds are chosen optimally like this? We have

$$PE(N) = \overbrace{\frac{1}{2(N+1)^2} + ... + \frac{1}{2(N+1)^2}}^{N+1 \text{ terms}} = \frac{1}{2(N+1)} \to 0, \text{ as } N \to \infty.$$

It is thus clearly advantageous to have as many thresholds as possible, i.e., to be able to perceive the world as finely as possible. Unfortunately, the ability to measure the world more precisely is biologically costly. Suppose the individual incurs a cost that is proportional to the probability of error as well as a cost $c(N)$ that depends directly on N, so that more thresholds are more costly. The total cost is then

$$PE(N) + c(N),$$

which should be minimized over the choice of N. If $c(N) \to 0$, in an appropriate uniform sense, it follows readily that $N \to \infty$ and $PE(N) \to 0$. As costs decline, the resulting choice behavior is exactly as conventionally predicted.

This exercise gives us some quite useful insights into how evolution would design a utility function to cope with a particular decision problem. One of the seemingly obvious but important lessons is that the optimal utility function depends upon the characteristics of the problem, in this case captured by the distribution F. Suppose evolution has to cope with different decision problems—sometimes one specification of F, sometimes another. Evolution would then like to tailor the utility function to each such problem, just as a different specification of F in our first example would give rise to a different utility function. To do so, however, evolution needs to "know" what problem the agent is facing.

This leads naturally to the second feature we seek in our analysis of adaptive utility and relative consumption effects, namely the relative consumption effects. The agent's past consumption or the consumption of others provides clues about the agent's decision environment and the choices the agent is likely to face. Evolution uses these clues to adjust the agent's utility, giving rise to a utility function that conditions current utilities on past consumption.

In examining this process, we follow Rayo and Becker [114]. Their model gives rise to two effects, namely,

(1) habituation—utility adjusts so that people get used to a permanent shift, positive or negative, in their circumstances, and

(2) peer comparisons—people are concerned with relative income or wealth.

What these have in common is a specification of utility in terms of a reference point that is determined either by one's own past consumption, or by the past and present consumption of peers. These are the relative consumption effects.

Rayo and Becker [114] again view utility as hedonic, as a biological device that induces appropriate actions by an individual. In particular, evolution chooses the mapping from material outcomes into pleasure in the most effective way possible. In the present context, this most effective way involves the construction of a reference point that reflects the individual's expectations of the world. As in Robson [122], there is a metaphorical principal-agent problem here, with evolution as the principal and the individual as the agent.

Evolution "wishes" the individual to be maximally fit, and she has the ability to choose the utility function of the agent to her best advantage. The key ingredients of the model are a limited range of utility levels that are possible, and a limited ability to make fine distinctions.[56]

Consider an agent who must choose a strategy $x \in X$. This might be interpreted as a method of hunting, for example, or more generally the pursuit of consumption. Once x is chosen, an output y is determined, with

$$y = f(x) + s$$

where the strictly concave function f represents the technology that converts the agent's consumption into output, and s is the realization of a random variable $\tilde{s}$ that has a zero mean and a continuous, unimodal density g, with $g' = 0$ only at its maximum. The agent must choose x before knowing the realization of $\tilde{s}$.

Evolution designs a utility function $V(y)$, attaching utilities to outputs, with the goal of maximizing the expected value of y. Notice that several familiar elements appear in this problem. First, evolution chooses a utility function to motivate the agent, rather than simply specifying or hard-wiring the optimal choice of x. The latter option is either prohibitively difficult, compared to the trial-and-error capabilities of evolution, or rendered impossible by a tendency for the technology f to change at a pace too rapid for evolution to supply corresponding adjustments in her prescription of x.[57] Second, while evolution's goal is the maximization of offspring, the variable y may represent directly observable intermediate goods such as money or food. Evolution then attaches utilities to values of y to induce choices that in turn have the desired effects in terms of offspring.

---

[56] Robson [121] argues that utility bounds and limited discrimination between utilities will induce evolution to induce adaptive utility functions that strategically position the steep part of the utility function. Tremblay and Schultz [153] provide evidence that the neural system encodes relative rather than absolute preferences, as might be expected under limited discrimination. See Friedman [52] for an early contribution and Netzer [97] and Wolpert and Leslie [163] for more recent work.

[57] We could capture this assumption more explicitly by writing the technology as $f(x, z)$, as do Rayo and Becker, where z represents features of the environment that affect the technology available to the agent and hence the agent's optimal actions, while assuming that the agent observes z but the possible values of z are too many and too complex for evolution to incorporate in the agent's utility function. Although the maximizer x then varies with the state z, the simplest Rayo and Becker formulation assumes that the maximized value of f does not. As we discuss briefly below, relaxing this assumption generates "S-shaped" utility functions rather than the step function derived for the simplest case. We omit z here in order to simplify the notation.

The agent's objective is to maximize

$$E\left\{V \mid x\right\} = \int V(f(x)+s)g(s)ds$$

over the choice of $x \in X$.

The first important constraint in the model is that there are bounds on V so that

$$V \in [\underline{V}, \overline{V}],$$

which we can then normalize so that $V \in [0,1]$. The constraints might ultimately reflect the fact that there are a finite number of neurons in the brain, and hence limits on the positive and negative sensations evolution can engineer the agent to produce. These upper and lower constraints on V will typically be binding, in that evolution would benefit from a wider range of emotional responses. It is expensive, however, to enlarge the range, and so this range must be finite and evolution must use the range optimally.

The second constraint is that the agent has only limited discrimination in distinguishing utilities. This takes the precise form that, if

$$\left|E\left\{V \mid x_1\right\} - E\left\{V \mid x_2\right\}\right| \leq \varepsilon,$$

then the individual cannot rank x_1 and x_2. Hence all choices within ε of $\max_{x\in X} E\left\{V \mid x\right\}$ are "optimal." It is assumed that the agent randomizes uniformly, or at least uses a continuous distribution with full support, over this satisficing set. Of course evolution would also prefer a smaller value of ε, but this is again expensive, and she will have to optimize given the optimal $\varepsilon > 0$.

Let x^* maximize $f(x)$. Then the agent thus chooses a value x from a satisficing set $[\underline{x}, \overline{x}]$, where

$$E\left\{V \mid x^*\right\} - E\left\{V \mid \underline{x}\right\} = E\left\{V \mid x^*\right\} - E\left\{V \mid \overline{x}\right\} = \varepsilon.$$

Evolution's goal is then to minimize the size of this satisficing set. The first step toward solving this problem is to note that evolution will maximize the difference in utilities between the optimal choice and the choice that lies just on the boundary of the satisficing set:

Lemma 7 *If V^* minimizes the satisficing set $[\underline{x}, \overline{x}]$, then V^* solves*

$$\max_{V(\cdot)\in[0,1]} E\left\{V \mid x^*\right\} - E\left\{V \mid \underline{x}\right\} \tag{19}$$

or, equivalently,

$$\max_{V(\cdot)\in[0,1]} = E\left\{V \mid x^*\right\} - E\left\{V \mid \overline{x}\right\}.$$

To verify this claim, suppose that it is not the case. Then, given the candidate optimum V^* and the attendant satisficing set $[\underline{x}, \overline{x}]$, there exists some other utility function $V \neq V^*$ such that

$$E\{V \mid x^*\} - E\{V \mid \underline{x}\} > E\{V^* \mid x^*\} - E\{V^* \mid \underline{x}\} = \varepsilon,$$

with, of course, an analogous inequality for $\overline{x}$. But then the alternative utility function V would give a smaller satisficing set, yielding a contradiction. This gives the result, and in the process a simple characterization of evolution's utility design problem.

It is now relatively straightforward to characterize the optimal utility function:

Proposition 8 *There exists a value $\hat{y}$ such that the optimal utility function V^* is given by*

$$V^*(y) = \begin{cases} 1 & y \geq \hat{y} \\ 0 & y < \hat{y} \end{cases}$$

where $\hat{y}$ solves

$$g(\hat{y} - f(x^*)) = g(\hat{y} - f(\underline{x})) = g(\hat{y} - f(\overline{x})).$$

To establish this, we recall that evolution's optimal utility function must minimize the satisficing set, which in turn implies that it must maximize the difference $E\{V \mid x^*\} - E\{V \mid \underline{x}\}$ (cf. (19). Writing the expectations in (19) and then changing variables to obtain the right side of the following equality, the utility function must be chosen to maximize

$$\int [V(f(x^*)+s) - V(f(\underline{x})+s)]g(s)ds = \int V(y)[g(y-f(x^*)) - g(y-f(\underline{x}))]dy.$$

Now the solution is clear. The smallest possible values of utility, or 0, should be assigned to values of y for which $g(y - f(x^*)) - g(y - f(\underline{x})) < 0$ and the largest possible utility, or 1, assigned to values of y for which $g(y - f(x^*)) - g(y - f(\underline{x})) > 0$. Our assumptions on g ensure that it has a "single-crossing" property, meaning that (since $f(x^*) > f(\underline{x})$) there is a value $\hat{y}$ that that $g(y - f(x^*)) - g(y - f(\underline{x})) < 0$ for all smaller values of y and $g(y - f(x^*)) - g(y - f(\underline{x})) > 0$ for all larger values. This gives the result. Notice that we could just as well have used $\overline{x}$ throughout this argument.

Evolution thus designs the agent with a "bang-bang" utility function, choosing a cutoff $\hat{y}$ such that outcomes above this cutoff induce the maximum possible utility, while those below minimize utility. As $\varepsilon \to 0$, the satisficing set collapses around x^* and the value of $\hat{y}$ approaches $f(x^*)$. Evolution thus becomes arbitrarily precise in penalizing the agent for choosing suboptimal values of x^*, as we would expect as the agent's perceptual imprecision disappears.

What lies behind this result? As a result of the agent's perceptual errors, evolution would like the utility function to be as steep as possible, so that the agent is routinely choosing between alternatives with large utility differences and hence making few mistakes. However, the constraints $\underline{V}$ and $\overline{V}$ on utility make it impossible to make the utility function arbitrarily steep everywhere. Evolution responds by making the utility function steep "where it counts," meaning over the range of decisions the agent is likely to encounter, while making it relatively flat elsewhere so as to squeeze the function into the utility bounds.

In the simple model presented here, making the utility function steep where it counts takes the extreme form of a single jump in utility. More generally, one might expect a smoother, S-shaped utility function to be more realistic than the cliff shape or bang-bang utility function we have derived. Notice first that the expected utility $E\{V|x\}$ that guides the agent's decisions has such an S shape. In addition, Rayo and Becker [114] show that an S shape would arise if deviations from a given reference level V_0 were costly. Alternatively, it might be that the agent knows more about the output technology than does evolution. Now evolution might not be able to target $E\{y|x^*\}$, instead having to smooth out V to provide strong incentives over a range of possible $E\{y|x^*\}$'s.[58]

Where do we see relative considerations in this model? We have the obvious beginnings of relative consumption effects in the need for evolution to tailor the utility function to the problem the agent faces, in order to position the "steep spot" at the appropriate place. Now suppose that output is given by

$$y = f(x) + s + w,$$

where w is a random variable whose value is observed by the agent before he makes his choice but is not observed by evolution, and s is again drawn subsequently to the agent's choice. The random variable w may capture aspects of the agent's environment that make high output more or less likely, while s captures idiosyncratic elements of chance and luck that affect the agent's output. Then evolution will condition the utility function on any

---

[58] Footnote 57 raised the possibility of incorporating an environmental variable z into the agent's technology, which would then be $f(x, z)$. As long as z affects only the shape of f, and hence the identify of the maximizer x^*, but not the value of the maximum $f(x^*, z)$, our previous analysis goes through without change. If z also affects the maximum $f(x^*, z)$, then the result is a smoother specification of the optimal utility function.

variables that carry information about w. If the agent is involved in a sequence of choices and there is persistence in the value of w, then evolution will condition the agent's current utility function on past realizations of the agent's output. A higher previous output will mean that it takes a higher current output to hit a given utility level. If the agent can observe others who are also affected by w, then evolution will condition the agent's utility function on the output of others. Observing higher output from one's neighbors will mean that a higher output must be produced to hit a given utility level. Relative consumption effects thus become the rule. Without such effects, trends in the value of w could eventually render the utility function irrelevant for the environment, with most choice occurring in a range where the utility function is virtually flat. All decisions would look equally good or bad and the individual's incentives would disappear.

For example, Rayo and Becker present a case in which $\hat{y}_t = y_{t-1}$. Hence, the individual is happy if and only if current output exceeds last period's output. Notice that in this case, the agent is punished as severely for bad luck as she would be for a bad decision. In equilibrium, the agent's decisions would be inevitably optimal and happiness would be purely a matter of luck.

This gives us the second of our desired features, namely a utility function that adjusts to reflect relative consumption effects. Finally, we can ask whether agents will anticipate these future adjustments when making their current choices, or will they remain unaware of such changes. Equivalently, will the agents be sophisticated or naive (cf. O'Donoghue and Rabin [100]). Robson and Samuelson [126] argue that evolution will optimally design agents to be at least partially naive. The intuition is straightforward. Suppose agents make intertemporal choices. Evolution then has conflicting goals in designing future utilities. On the one hand, they must be set so as to create the appropriate tradeoffs between current and future consumption, so that agents have appropriate investment incentives. On the other hand, once the future is reached, evolution would like to adjust the utility function to create the most effective current incentives.

These forces potentially conflict. Suppose that current investment can create lucrative future payoffs. Evolution would like to promise high future utilities, in order to induce such investment. Once the investment has been made and the future reached, however, evolution would like to ratchet the entire utility function down, so as to continue to create incentives. But an agent who anticipates this will not undertake the current investment. The solution? Make the agent naive, so that she has current investment incentives in anticipation of lucrative future payoffs, which are subsequently and unexpectedly adjusted so as to heighten subsequent incentives.

4.2.3 Implications

In each of the two preceding subsection, we find utility functions that that are defined over the consumption of others as well as one's own consumption, providing foundations for preferences that are not pure "selfish." In each case, these relative consumption effects implicity incorporate useful environmental information into the agent's utility maximization.

Why do we care about such relative consumption effects? What behavior might we expect to observe that is consistent with relative consumption effects? Why do we care whether they might enter preferences directly? We take these questions in reverse order.

Our current world is much different from the ancestral environment in which our preferences evolved. If we were concerned only with the ancestral environment, then our interest would not extend beyond the behavior that maximizes fitness. We would be interested in whether behavior exhibited relative consumption effects, but we could ignore imperfections such as the agent's noisy information processing that have only a minor impact (or, in the case of our simple model, no impact) on the constrained-optimal behavior implemented by evolution. If we are concerned with our current world, however, then we must recognize that these imperfections can have a important impact on the mechanism by which evolution induces her optimal behavior, and that the implementing mechanism can in turn have an important impact on the behavior that appears once the agents are transplanted from the ancestral environment to our much different modern environment. For example, perfect Bayesians will never erroneously imitate uninformative consumption decisions. Relative consumption effects that are embedded in preferences may cause agents in a modern environment to condition their behavior on a variety of uninformative or misleading signals, regardless of the uncertainty they face. It thus makes a difference what sort of behavior evolution has programmed us to have, and how she has done the programming.

What would we expect to see in a world of relative consumption effects? First, we should see evidence that evolution designs agents to either consciously or unconsciously make use of environmental cues in shaping consumption decisions. Experiments have shown that some animals condition their fat accumulation on day length, a source of information that is reasonably reliable in natural environments but that can be used to manipulate feeding behavior in laboratory settings (Mercer, Adam and Morgan [92]). A variety of young animals, including humans, have been shown to be more likely to consume foods that they have observed others consuming (Smith

[145, Section 2.1]). More striking is recent evidence that a low birth weight puts one relatively at risk for subsequent obesity (Petry and Hales [105], Ravelli, van der Meulen, Osmond, Barker and Bleker [112]). The conventional interpretation is that poor maternal nutrition is a prime contributor to a low birth weight as well as a prime indicator of a meager environment, so that a low birth weight provides information to which the optimal reaction is a tendency to store more bodily food reserves.

In addition, we should observe an inclination to conform to the behavior of others that will sometimes appear to be unjustified on informational grounds. Psychologists again commonly report a taste for conformity (Aronson [5, Chapter 2], Cialdini [24, Chapter 4]), even in situations in which one would be extremely hard-pressed to identify an objective information-based reason for doing so.[59]

Our model of relative consumption effects directs attention to conformity effects that initially appear somewhat counterintuitive. The model suggests that relatively low-productivity agents will strive to increase consumption, while high productivity agents will attenuate their consumption, both in order to not be too conspicuously different. The latter finding contrasts with the popular view of relative consumption effects as creating incessant incentives to consume more in order to "keep up with the Joneses." Do we expect internet billionaires to lie awake at night, desperately searching for ways to dispose of their wealth so as to look more like ordinary people? Notice first that information-based relative consumption effects are consistent with outcomes in which some people happily, even gloatingly, consume more than others, perhaps much more. Higher-productivity agents optimally consume more than lower-productivity agents, both in the model and in the world. The billionaire need not lie awake at night.

More importantly, the behavior predicted by the model is that agents who observe others consuming more should themselves consume more. But this is typically what one means by "keeping up with the Joneses." Information-based relative consumption effects imply not that we must observe people endeavoring to reduce their consumption, but rather that people whose characteristics lead to high consumption levels should strive less vigorously to keep ahead of the Joneses than they would to catch up if the Joneses were ahead.

Preferences incorporating relative consumption effects give rise to the

---

[59] The work of Asch [6] is classic, in which an apparent desire to conform prompted experimental subjects to make obviously incorrect choices when matching the lengths of lines, while denying that they were influenced by the choices of others.

risk that agents will react to others' consumption in ways that do not reflect the informational content of their surroundings, leading to outcomes that are inefficient (conditional on the environment). Evolution may have optimally incorporated these risks in the ancestral environment in which our preferences evolved, but new problems appear as agents apply their behavioral rules to a modern industrial society for which they are likely to be a poor match.[60] In addition, to the extent that evolution has responded to this risk, she has done so to maximize the fitness of her agents. From our point of view, it is utility and not fitness that counts. Studying evolutionary foundations allows us to gain insight into the difference between evolution's preferences in the ancestral environment and our preferences in our current world, in turn helping us assess modern social developments or policy interventions.

For example, it is likely that the observations which motivate information-based relative consumption effects are stratified, with evolution finding it optimal for her agents to react more strongly to the generally more relevant consumption of others who appear to be "like them" than to people whose circumstances are quite different. Hence, we may be unfazed by comparisons with internet billionaires, but may be much more conscious of how our consumption compares with that of our colleagues. However, the concept of likeness on which such stratification is based is likely to be both endogenous and liable to manipulation. The development of modern advertising and mass communications may accentuate the visibility of high consumption levels and hence the inefficiencies caused by relative consumption effects. Information and communication technologies may thus bear a hidden cost.

Suppose next that we consider an inequality policy designed to decrease the variation in individual productivities, perhaps by enhancing the productivity of those at the bottom of the income and consumption scale. This will tend to compress the distribution of consumption levels. Consumers will thus observe others who look more like themselves, attenuating the distortions caused by information-based relative income effects. In contrast, if agents seek status that is tied to conspicuous consumption, then compressing the distribution of consumption increases the returns to investing in status, since a given increase in consumption now allows one to "jump over" more of one's contemporaries. The result can be a ruinous race to invest in status, possibly making everyone worse off (Hopkins and Kornienko [74]). Policy prescriptions can thus depend critically on whether relative

---

[60] For example, Frank [47] argues that relative consumption effects lead Americans to undersave, overconsume luxury goods, and underconsume leisure and public goods.

consumption effects arise out of information or status concerns.

4.3 Group Selection

Much of the recent interest in more sophisticated models of preferences has been motivated by the belief that people are not as relentlessly selfish as economic models might have us believe. People donate to charity, they vote, they provide public goods, they come to the aid of others, and they frequently avoid taking advantage of others. Such "other-regarding" behavior is often invoked as one of the distinguishing and puzzling features of human society (e.g., Seabright [139]). At first glance, however, evolutionary arguments appear particularly unlikely to generate other-regarding behavior. Where else would the survival of the fittest lead but to relentless self-interested behavior? Upon closer reflection, there is ample room for evolution to generate more complex and other-regarding preferences. Perhaps the leading candidate for doing so is the familiar concept of group selection, by which evolution can seemingly design individuals whose behavior is beneficial to the group to which they belong. It is accordingly only natural that we touch here on the idea of group selection.

It is uncanny how close Darwin came to the modern view of biological evolution, given that a detailed understanding of the mechanics of genetic inheritance lay far in the future. In particular, he emphasized that a certain variation would spread if this variation led to greater reproductive success for *individuals* and was inherited by their descendants. We now have a better understanding of the genetics behind the inheritance, as well as a more nuanced view of whether it is the individual, the gene, or something else that is the appropriate unit of selection, but the basic understanding remains the same.

At the same time, Darwin occasionally wandered away from models of evolution based in the fates of individuals, into what would now be called "group selection." Thus, he thought an individual human might engage in behavior that is beneficial to the survival of a group, even if this behavior had a fitness cost to the individual. To what extent can group selection help us explain our preferences?[61]

There is a "folk wisdom" appeal to group selection, and this mechanism was once routinely invoked in popular accounts of natural selection. For example, the idea that a predator species was doing a prey species a favor by eliminating its weakest members represented one of the more fanciful

---

[61] This section is based on Robson [123].

extremes in applying "group selection" arguments. More scientifically, the English experimental biologist Wynne-Edwards [164, 165] opened the modern discussion of group selection by providing a clear and explicit manifesto on group selection, in the process becoming a favorite target for those wishing to preserve a focus on the individual (or gene). For example, he argued that birds limit the size of their clutches of eggs to ensure that the size of the population does not exceed the comfortable carrying capacity of the environment. That is, individuals act in the best interest of the species, with those that do so most effectively being evolutionarily rewarded by the resulting success of their species.

These early group selection arguments were effectively devastated by Williams [160]. If a new type of individual does not so obligingly limit her clutch, for example, why would this more fertile type not take over the population, even though the result is disastrous for the population's standard of living? After all, the profligate egg-layer inevitably has more offspring than her more restrained counterparts, even if the result is counterproductive overcrowding. This challenge to the logic of group selection was complemented by doubts as to the need for group selection. For example, one can find compelling arguments as to why it is in the interests of an individual to limit her clutch size. It might be that, beyond a certain point, an increase in the number of eggs reduces the expected number of offspring surviving to maturity, because each egg then commands a reduced share in parental resources. A finite optimum for clutch size is then to be expected. Thus, observations suggesting that clutch sizes are limited do not compel a group selection interpretation. As a collection of similar observations accumulated, some biologists were tempted to argue that evolutionary theory could dispense with group selection entirely. Dawkins [32] has been especially insistent in rejecting group selection, in the process going further in the other direction by arguing for the primacy of the gene rather than individual as a still more basic unit of selection.

Subsequent work suggests that there certainly are phenomena best understood at the level of the gene, but at the same time has uncovered cases in which evolution appears to proceed at different levels. Consider, for example, meiotic drive, also known as segregation distortion. This refers to any process which causes one gametic type to be over-represented or under-represented in the gametes formed during meiosis, and hence in the next generation. A classic example of meiotic drive concerns the T locus in mice. This locus controls tail length, but also the viability of the mouse. The following facts apply—TT homozygotes have normal long tails, Tt heterozygotes have short tails, which is presumably somewhat disadvantageous, and *tt* homozygotes are sterile. If this were the whole story, there would be unambiguous selection against the *t* allele. However, the wrinkle is that the *Tt* heterozygotes transmit the *t* allele with about probability 90% to their sperm, rather than the usual Mendelian 50%. Hence when the *t* alelle is rare, this strong meiotic drive will overcome the slight fitness disadvantage of short tails and the frequency of the *t* allele will increase. Eventually, the *tt* homozygotes will occur with appreciable frequency, and there will be an equilibrium mixture of the two alleles. The evolutionary processes governing tail length in mice thus mixes considerations that arise at two levels of selection: positive selection for *t* haplotypes at the level of the gene, but negative selection for *tt* individuals at the level of the organism. But if selection can operate at both the genetic and individual level, might it not sometime also operate at the group level?

We want to be clear in recognizing the primacy of the gene as the unit of evolutionary selection. It is genes that carry characteristics from one generation to the next, and only through genes can characteristics be inherited. At the same time, genes are carried by individuals, and which genes are relatively plentiful can depend upon the fate of their host individuals. But could not the fate of these individuals depend upon the groups to which they belong?

We address these issues by examining the interplay between individual and group selection. Again, we emphasize the importance of beginning with the perspective of the gene. However, there are many cases where the interests of the gene and the individual do not conflict. In addition, it is often difficult to give concrete form to the notion of the gene as the unit of selection, given our ignorance of the details of the transformation of genes into individual traits, particularly for complex behavioral characteristics.[62] Hence, despite the theoretical primacy of the gene, we believe we can usefully simplify the analysis by restricting attention here to the comparison between individual level and the group level of selection.

4.3.1 The Haystack Model

In order to fix ideas, we consider the classic haystack model, offered by Maynard Smith [90] to study the issue of individual selection versus group selection. Our account simplifies the standard model in several ways. Perhaps most importantly, reproduction here is assumed here to be asexual.

---

[62] Grafen [62] advocates finessing such detailed questions on the genetic basis of individual variation, an argument refereed to as his "phenotypic gambit."

There are a number of haystacks in a farmer's field, where each haystack is home to two mice. Each pair of mice plays the prisoners' dilemma, choosing between the usual two alternatives—cooperate or defect—and then dies. However, each individual leaves behind a number of offspring equal to her payoff in the prisoners' dilemma. The heritable characteristic of an individual is her choice to either cooperate or defect, so we can think of the population as being divided between cooperators and defectors. In particular, offspring inherit their mother's choice of strategy.

After this initial play of the prisoners' dilemma by the haystack's founding pair, there are a number $T-1$ of subsequent stages of play, where the mice in each haystack are paired at random, play the prisoners' dilemma, and then die, while giving rise to further offspring in numbers determined by their prisoners'-dilemma payoffs. The number of individuals within the haystack choosing each strategy then grows in an endogenous fashion, as does the overall size of the group. Every so often, once a year, say, the haystacks are removed, and the mice are pooled into a single large population. Now pairs of mice are selected at random from the overall population to recolonize the next set of haystacks, and excess mice die.

To give an example, consider the following version of the prisoners' dilemma:

$$\begin{array}{c|c|c|} & C & D \\ \cline{2-3} C & 2,2 & 0,4 \\ \cline{2-3} D & 4,0 & 1,1 \\ \cline{2-3} \end{array}\ .$$

As a further simplification, suppose that there are a large number of haystacks and therefore individuals, although this assumption facilitates group selection and hence is not innocent. Suppose that the initial fraction of C's in the population is $f \in [0,1]$. Hence the fraction of haystacks that are colonized by 2 C's is f^2; the fraction that are colonized by 2 D's is $(1-f)^2$; and the fraction that have one of each is $2f(1-f)$. There are T rounds of play within each haystack. It follows that each pair of C's gives rise to 2^{T+1} descendants, who are also C's. Each pair of D's gives rise to just 2 D's. Each pair of one C and one D gives rise to 4 D's.

At the end of the T periods of play, and hence just as the haystacks are disrupted, the new fraction of C's in the population is,

$$f' = \frac{2^{T+1}f^2}{2^{T+1}f^2 + 8f(1-f) + 2(1-f)^2}. \tag{20}$$

Let us check first what happens if $T = 1$. In this case, $f' < f$ if and only if

$$4f < 4f^2 + 2(1-f)(3f+1) = 2 + 4f - 2f^2 \quad \Leftrightarrow \quad f < 1.$$

That is, in this case, the D's will increase, and $f \to 0$. This is not surprising, since with $T = 1$, we simply have an elaborate description of the usual prisoners' dilemma—the extra generality available in the structure of the haystack model is not used. Pairs are broken up immediately so that there is no opportunity to exploit the relatively high total payoffs for the haystack/group that arise from two initial C's.

When there is more than one generation per haystack cycle, these relatively high total payoffs may quickly outstrip those from any other possible starting combination of mice. In particular, if $T \geq 3$, then we have $f' > f$ as long as f is close enough to 1. To see this, we use (20) to conclude that more cooperators than defectors will emerge from the haystacks if

$$2^{T+1}f > 2^{T+1}f^2 + 8f(1-f) + 2(1-f)^2 = 2^{T+1}f^2 + 2(1-f)(3f+1)$$

which in turn holds if

$$T(f) = 2^{T+1}f^2 + 2(1-f)(3f+1) - 2^{T+1}f < 0.$$

Moreover, there is some $f < 1$ sufficiently large as to satisfy this inequality for all $T \geq 3$, an observation that follows immediately from noting that

$$T(1) = 0, \text{ and } T'(1) = 2^{T+2} - 8 - 2^{T+1} = 2^{T+1} - 8 > 0.$$

Hence, in this case, the relatively high growth rate of groups founded by cooperators is sufficiently strong as to allow cooperation to dominate a population whose initial proportion of cooperators is sufficiently large. Cooperation is rescued in the prisoners' dilemma by group selection.

Maynard Smith's intention in examining this model was to give the devil his due by identifying circumstances under which group selection might well have an effect. At the same time, he regarded the analysis as making it clear that the assumptions needed to make group selection comparable in strength to individual selection would be unpalatable. First, in order for group selection to be effective in the haystack model, there must obviously be a number of groups, preferably a large number.

Second, there must be a mechanism that insulates the groups from one another. Only then can a cooperative group be immune to infection by a defecting individual, and hence be assured of maintaining its greater growth rate. Groups must thus be isolated from the appearance of migrating Ds as well as D mutants. Third, even with the temporary insulation of each haystack in this model, cooperation will only evolve if there are sufficient rounds of play within each haystack, so that cooperation amasses a sufficient advantage as to survive the next sampling.

While there is some room to relax these assumptions, and one might hope that alternative models are more amenable to group selection, a reasonably widespread view within biology is that group selection is logically coherent but of limited importance.[63] The requirements of a large number of groups, sufficient isolation of groups, barriers to migration and mutation, and differential group success rates, all combine to limit the applicability of group selection. Intuitively, a loose description of the problem with group selection is that it relies too heavily upon the assumption that a bad choice will lead to *group* extinction. There is clearly scope in reality for individual selection, since individuals die frequently, but the idea that groups face extinction sufficiently often as to potentially overwhelm the strength of individual selection strikes many as less plausible.

4.3.2 Selection Among Equilibria

Much of the initial attention was devoted to the possibility of group selection leading to different results than would individual selection, as in the prisoners' dilemma, in the process leaving many skeptics as to the effectiveness and importance of group selection. However, there is a compelling alternative scenario in which group selection may well operate robustly, in any species. This is as a mechanism to select among equilibria (Boyd and Richerson [16, 17]).

Consider a population that is divided into various subpopulations, largely segregated from one another, so that migration between subpopulations is limited. The members of each subpopulation are randomly matched to play the same symmetric game, which has several symmetric equilibria. For example, suppose the game is the simplest 2×2 coordination game:

$$\begin{array}{c|c|c|} & A & B \\ \hline A & 2,2 & 0,0 \\ \hline B & 0,0 & 1,1 \\ \hline \end{array} \quad .$$

Individual selection ensures that some equilibrium is attained within each subpopulation. In general, some subpopulations would play the A equilibrium, and some would play the B equilibrium. Each of these configurations is internally robust. That is, if there were the occasional B arising by mutation in an A subpopulation, it would find itself at a disadvantage and would die out. Similarly an A mutant in a B population would die out, despite the ultimate advantage of getting to the all-A configuration. Alternatively,

---

[63] See Sober and Wilson [146] for a forcefully argued alternative view.

a small group of individuals may occasionally migrate from one subpopulation to another. If the newcomers did not match the prevailing action in their new subpopulation, the newcomers will once again disappear.

Now consider the competition between subpopulations. The A subpopulations grow faster than do those that play B. It is then reasonable to suppose the B populations will eventually die out completely. That is, group selection is free to operate in a leisurely fashion to select the Pareto superior equilibrium. There is no tension here between the two levels of selection, and hence no calculations that need to be made about the number of groups or rates of mutation and migration. Indeed, given enough time, virtually any group structure will lead to a population dominated by the Pareto superior equilibrium. The implication, in Boyd and Richerson's [16, 17] view, is that group selection theories have missed the boat by concentrating on the prisoners' dilemma. The true strength of group selection may be not to motivate behavior at odds with individual selection, but as a force mitigating between various contenders for the outcome of individual selection.

4.3.3 Group Selection and Economics

Why does group selection matter in economics? Group selection is one of the most obvious mechanisms for generating preferences in humans to behave in the social interest rather than that of the individual. At stake then is nothing less than the basic nature of human beings.

As an economist, one should be skeptical of the need to suppose that individuals are motivated by the common good. Economic theory has done well in explaining a wide range of phenomena on the basis of selfish preferences, and so the twin views of the individual as the unit of selection and as the extent of the considerations that enter one's utility function are highly congenial to economists. Furthermore, to the extent that armchair empiricism suggests that non-selfish motivations are sometimes present, these seem as likely to involve malice as to involve altruism. For example, humans seem sometimes motivated by relative economic outcomes, which apparently involves a negative concern for others. Finally, group selection is a potentially blunt instrument that might easily "explain" more than is true.

There are, nevertheless, some aspects of human economic behavior that one is tempted to explain by group selection. For example, human beings are often willing to trade with strangers they will likely never see again, behavior that might be analogous to cooperating in the one-shot prisoners' dilemma. Indeed, there is no shortage of reliable data showing that human beings are capable of such apparently irrationally cooperative behavior, in appropriate circumstances. Whatever the underlying reasons for this, it is a significant factor in supporting our modern economic and social structure.

One possibility is that we are simply mistaken in likening this behavior to cooperation in the prisoners' dilemma. It might be that we trade with others rather than simply trying to seize their goods because there are effective sanctions for behaving otherwise. Alternatively, it is sometimes argued that the structure of the hunter-gatherer societies characteristic of our evolutionary past helps account for cooperative behavior in modern settings. Hunter-gatherer societies were composed of a large number of relatively small groups, and individuals within each group were often genetically related. Perhaps, so the argument goes, we acquired an inherited psychological inclination towards conditional cooperation in such a setting, partly perhaps as a result of group selection. The group selection argument here gets a boost not only from a setting in which small, relatively isolated groups is likely to have been the norm, but from the fact that the members of these groups were likely to be related, allowing group selection to free ride on the forces of kin selection.[64] The resulting cooperative inclinations may then have carried over into modern societies, despite genetic relatedness now being essentially zero on average.

It is hard to believe, however, that hunter-gatherers never encountered strangers, and that it wasn't important to both keep track of who was a stranger and to adjust one's behavior accordingly. If there were good reasons to condition on this distinction, why would corresponding different strategies not have evolved? Why wouldn't we now use the "defect against strangers" response nearly always? Even if we did somehow acquire a genetic inclination to cooperate in archaic societies, shouldn't we now be in the process of losing this inclination in modern large and anonymous societies?

Sober and Wilson [146] push energetically for a rehabilitation of group selection within biology. They argue that kin selection—the widely accepted notion that individuals are selected to favor their relatives—should be regarded as a special case of group selection. Proceeding further, they note that what matters most fundamentally is the likelihood that altruistic individuals will be preferentially matched with other altruistic individuals. They offer kin selection as one obvious circumstance under which this will be the case, while arguing that there are many others. While kin selection is widely accepted, one must remember that the mechanisms for achieving the preferential matching of altruistic individuals are quite different for kin

---

[64] See Eshel [41] for a discussion of the relationship between kin selection and group selection.

selection and group selection. In the end, a skeptical view of the importance of group selection appears to be common among biologists.

4.3.4 Implications

Of all the topics considered in this essay, group selection has perhaps the widest range of potential applications. With the appropriate model, group selection allows us to rationalize almost any behavior. This may explain why biologists, though readily conceding the logical coherence of group selection arguments, typically exhaust all other avenues before turning to group selection as an explanation.[65] We view finding ways to assess group selection arguments, and to separate those circumstances in which group selection is an essential element of an explanation from those in which it provides a convenient alternative story, as one of the foremost challenges facing those working on the evolutionary foundations of economic behavior.

5 Concluding Remark

This essay has addressed a broad subject area, and has all too predictably touched only a fraction of it, despite consuming many pages. We believe there is much to be learned, and much yet to be done, in studying the evolutionary foundations of economic behavior. Pursuing these topics should bring economists increasingly into contact with work in biology and psychology, both of which have much to offer. We have no doubt that we can continue to produce elegant evolutionary models. Will they remain simply nice models, or will they serve as the basis for the type of applied work that motivates our interest in them? This key question remains to be answered. An affirmative answer will require moving beyond the theoretical foundations with which this essay has been concerned to demonstrate that these models are useful in addressing particular applied questions. Can they help us get better estimates of patterns of risk aversion or discounting? Can they help us design more effective economic institutions? There is clearly much work to be done.

---

[65] One is reminded in this respect of Wilson's [162] caution to economists that reputation models may well make things too easy to explain.

6 Proofs

6.1 Proof of Proposition 1

We provide the proof for the case in which $N(0) = \left(\frac{1}{T}, \dots, \frac{1}{T}\right)$. Relaxing this assumption requires only more tedious notation.

Fix a time t. Let τ_t identify the event that the period-t Leslie matrix features $x_\tau \neq 0$ (and all other $x_{\tau'} = 0$). We say in this case that environment τ_t has been drawn in period t. Then only parents of age τ_t reproduce in period t, having x_{τ_t} offspring. There are $S^{\tau t} N_0(t - \tau_t)$ such parents, so that we have

$$N_0(t) = S^{\tau t} x_{\tau_t} N_0(t - \tau_t).$$

We can perform this operation again. Let $\tau_{t-\tau_t}$ be the environment drawn at time $t - \tau_t$. Then we have

$$N_0(t) = S^{\tau t} x_{\tau_t} S^{\tau t-\tau t} x_{\tau_{t-\tau_t}} N_0(t - \tau_t - \tau_{t-\tau_t}).$$

Continuing in this fashion, we have

$$N_0(t) = S^t x_{\tau_t} x_{\tau_{t-\tau_t}} x_{\tau_{t-\tau_t-\tau_{t-\tau_t}}} x_{\tau_{t-\tau_t-\tau_{t-\tau_t}-\tau_{t_{t-\tau_t-\tau t-\tau_t}}}} \cdots \frac{1}{T},$$

for a sequence $\tau_t, \tau_{t-\tau_t}, \tau_{t-\tau_t-\tau_{t-\tau_t}}, \tau_{t-\tau_t-\tau_{t-\tau_t}-\tau_{t_{t-\tau_t-\tau_{t-\tau_t}}}}, \dots$ with the property that τ_t is the environment drawn in period t, $\tau_{t-\tau_t}$ is the environment drawn in period $t - \tau_t$, $\tau_{t-\tau_t-\tau_{t-\tau_t}}$ is the environment drawn in period $t - \tau_t - \tau_{t-\tau_t}$, and $\tau_{t-\tau_t-\tau_{t-\tau_t}-\tau_{t_{t-\tau_t-\tau_{t-\tau_t}}}}$ is the environment drawn in period $t - \tau_t - \tau_{t-\tau_t} - \tau_{t_{t-\tau_t-\tau_{t-\tau_t}}}$, and so on. The $1/T$ represents the initial mass of parents of the appropriate age, and the sequence $\tau_t, \tau_{t-\tau_t}, \dots, \tau_{t'}, \tau_{t''}$ has the properties

$$\tau_t + \tau_{t-\tau_t} + \dots + \tau_{t'} \quad < \quad t \tag{21}$$

$$\tau_t + \tau_{t-\tau_t} + \dots + \tau_{t'} + \tau_{t''} \quad \geq \quad t. \tag{22}$$

Hence, the final environment in this sequence, $\tau_{t''}$, causes offspring to survive who are born to a generation of parents that were alive at time 0. The age of these parents at time 0 depends upon the period in which $\tau_{t''}$ is drawn and the realization of $\tau_{t''}$, and may be any of the generations alive at time 0. Since there are $1/T$ of each age at time 0, the final $1/T$ is applicable regardless of which time-0 age is relevant.

We can then write

$$N_0(t) = \frac{1}{T} S^t \prod_{\tau=1}^{T} x_\tau^{r_\tau(t)}$$

and hence, taking logs and then dividing by t,

$$\frac{1}{t}\ln N_0(t) = \ln S + \sum_{\tau=1}^{T} \frac{r_\tau(t)}{t}\ln x_\tau - \frac{\ln T}{t}, \tag{23}$$

where $r_\tau(t)$ is the number of times environment τ is drawn in the sequence $\tau_t, \tau_{t-\tau_t}, \tau_{t-\tau_t-\tau_{t-\tau_t}}, \tau_{t-\tau_t-\tau_{t-\tau_t}-\tau_{t_{t-\tau_t-\tau_{t-\tau_t}}}}, \ldots, \tau_{t''}$. Our analysis then rests on examining the numbers $r_1(t), \ldots, r_T(t)$. Notice that so far, we have made no use of independence assumptions, having only rearranged definitions. Independence plays a role in examining the $r_\tau(t)$.

Intuitively, the argument now proceeds along the following lines:

- As t gets large, each of the $r_\tau(t)/t$ converges to R_t/Tt, where R_t is the total number of draws in the sequence, i.e., the proportion of periods featuring a draw of environment τ is very nearly the same for all $\tau = 1, \ldots, T$. This follows from the observations that each environment is equally likely and environments are drawn independently each time one is drawn, and gives

$$\lim_{t\to\infty} \sum_{\tau=1}^{T} \frac{r_\tau(t)}{t}\ln x_\tau \quad = \quad \lim_{t\to\infty} \sum_{\tau=1}^{T} \frac{R_t}{Tt}\ln x_\tau .$$

- From (21)–(22), the total number of draws R_t is determined approximately (with the approximation arising out of the fact that the parents of those offspring who survive as a result of draw $\tau_{t''}$ may be older than 1 at the beginning of the process, and with the approximation thus becoming arbitrarily precise as the number of draws increases) by

$$\sum_{\tau=1}^{Tt} \frac{R_t}{T}\tau = \frac{R_t}{Tt}\sum_{\tau=1}^{T} \tau = 1.$$

- This is the statement that the total of the reproductive lengths drawn in the course of the sequence $\tau_t, \tau_{t-\tau_t}, \tau_{t-\tau_t-\tau_{t-\tau_t}}, \tau_{t-\tau_t-\tau_{t-\tau_t}-\tau_{t_{t-\tau_t-\tau_{t-\tau_t}}}}$, $\ldots, \tau_{t''}$ must equal t. This gives

$$\lim_{t\to\infty} \sum_{\tau=1}^{T} \frac{r_\tau(t)}{t}\ln x_\tau \quad = \quad \frac{\sum_{\tau=1}^{T} \ln x_\tau}{\sum_{\tau=0}^{T} \tau}.$$

  Inserting this in (23) gives (8), the desired result.

Our first step in making this argument precise is to confirm that the random draws determining the environments in the sequence $\tau_t, \tau_{t-\tau_t}, \tau_{t-\tau_t-\tau_{t-\tau_t}}, \tau_{t-\tau_t-\tau_{t-\tau_t}-\tau_{t_{t-\tau_t-\tau_{t-\tau_t}}}}, \ldots, \tau_{t''}$ are independent. This is not completely obvious. While the environment is determined independently in each period, the identities of the periods at which the draws are taken in this sequence are endogenously (and hence randomly) determined, potentially vitiating independence.

To examine this question, we construct a model of the stochastic process determining the environment. Consider the measure space $([0,1], \mathcal{B}, \lambda)$, where λ is Lebesgue measure and $\mathcal{B}$ is the Borel σ-algebra. We now model the process determining the environment by letting $\xi(1)$ be a random variable defined by

$$\omega \in \left(\frac{\tau-1}{T}, \frac{\tau}{T}\right) \Rightarrow \xi(1)(\omega) = \tau, \quad \tau = 1, \ldots, T.$$

We then define $\xi(2)$ by

$$\omega \in \left\{\left(h + \frac{\tau-1}{T^2}, h + \frac{\tau}{T^2}\right) \text{ for some } h \in \{0, 1, \ldots, T\}\right\} \Rightarrow \xi(2)(\omega) = \tau, \quad \tau = 1, \ldots, T.$$

Continuing in this fashion gives a countable sequence of random variables that are independent and that each are equally likely to take each of the values $1, 2, \ldots, T$. We interpret $\xi(t)$ as determining the environment at time t. But it is now a straightforward calculation that

$$\Pr\{\xi(t) = \tau, \xi(t-i) = \tau'\} = \frac{1}{T^2}$$

for any τ and τ', and hence that $\xi(t)$ and $\xi(t-\tau_t)$ are independent. This in turn ensures that the sequence $\tau_t, \tau_{t-\tau_t}, \tau_{t-\tau_t-\tau_{t-\tau_t}}, \tau_{t-\tau_t-\tau_{t-\tau_t}-\tau_{t_{t-\tau_t-\tau_{t-\tau_t}}}}, \ldots, \tau_{t''}$ is independent.

Let

$$K \equiv \sum_{\tau=1}^{T} \tau.$$

Our goal is to show that with probability one,

$$\lim_{t\to\infty} \frac{r_\tau(t)}{t} = \frac{1}{K}, \tag{A5}$$

which combines with (26) to imply (15), giving the desired result.

We now construct a model of the process determining the frequencies $r_\tau(t)$. To do this, consider again the measure space $([0,1], \mathcal{B}, \lambda)$, where λ is Lebesgue measure and $\mathcal{B}$ is the Borel σ-algebra. Let $\zeta(1)$ be a random variable defined by

$$\omega \in \left(\frac{\tau-1}{T}, \frac{\tau}{T}\right) \Rightarrow \zeta(1)(\omega) = \tau, \quad \tau = 1, \ldots, T.$$

We then define $\zeta(2)$ by

$$\omega \in \left\{\left(h + \frac{\tau-1}{T^2}, h + \frac{\tau}{T^2}\right) \text{ for some } h \in \{0,1,\ldots,T\}\right\} \Rightarrow \zeta(2)(\omega) = \tau, \ \ \tau = 1, \ldots, T.$$

Continuing in this fashion again gives a countable sequence of random variables that are independent and that each are equally likely to take each of the values $1, 2, \ldots, T$. In particular, having fixed t, we think of $\zeta(1)$ as describing the draw of the environment at time t. Then, noting that $\zeta(2)$ is independent of $\zeta(1)$ and has the same distribution as $\xi(t - \tau_t)$ regardless of the value of τ_t, we think of ζ_2 as describing the draw of the environment at time $t - \tau_t$. Similarly, $\zeta(3)$ describes the draw at time $t - \tau_t - \tau_{t-\tau_t}$, and so on. The frequencies $r_\tau(t)$ thus are determined by the draws from the collection $\zeta(1), \ldots, \zeta(\hat{t}(t))$ for some number $\hat{t}(t)$. The time $\hat{t}(t)$ is randomly determined and is given by

$$\hat{t}(t) = \max\{t : \sum_{s=0}^{t-1} \tau_s < t\}. \tag{A6}$$

Then $r_\tau(t)$ is the number of times environment τ is drawn by the random variables $\zeta(1), \ldots, \zeta(\hat{t}(t))$.

Fix $\epsilon > 0$ and define $t'(t)$ (hereafter typically written simply as t') to satisfy

$$t'(t)\left(\left(\frac{1}{T} - \epsilon\right)K + T^2\epsilon\right) = t. \tag{A7}$$

Notice that $t > t'(t)$ (this is equivalent to $T^2 > K$) and that t' is linear and increasing in t. Intuitively, $t'(t)$ will be useful because (as we will see) with high probability $t'(t) < \hat{t}(t)$, i.e., with high probability, the random stopping time has not yet been encountered by time $t'(t)$.

Let $\rho_i(t')$ be the number of times environment i is drawn by the random variables $\zeta(1), \ldots, \zeta(t')$. Then choose t and hence $t'(t)$ sufficiently large that, with probability at least $1 - \epsilon$, we have

$$\frac{1}{T} - \epsilon < \frac{\rho_\tau(t')}{t'} < \frac{1}{T} + \epsilon \tag{A8}$$

for $\tau = 1, \ldots, T$. The weak law of large numbers ensures the existence of such t. Let $\Sigma \subset [0, 1]$ be the event that these inequalities hold (and note that $\lambda(\Sigma) \geq 1 - \epsilon$). For our purposes, the key characteristic of Σ is that on Σ,

$$t' \left(\left(\frac{1}{T} - \epsilon \right) K + T\epsilon \right) \leq \sum_{s=1}^{t'} \zeta(s) \leq t' \left(\left(\frac{1}{T} - \epsilon \right) K + T^2\epsilon \right) = t. \qquad \text{(A9)}$$

The term $\sum_{s=1}^{t'} \zeta(s)$ is the sum of the realizations of the t' random variables $\zeta(1), \ldots, \zeta(t')$. The left term is the smallest value this sum can take on Σ, which is obtained by first assuming that every value $i \in \{1, \ldots, T\}$ appears just often enough to attain the minimum frequency $\frac{1}{T} - \epsilon$ (giving the term $\left(\frac{1}{T} - \epsilon\right) K$), and then that all additional draws ($t'(1 - (\frac{1}{T} - \epsilon)T) = t'T\epsilon$ of them) all give environment 1. The third term is the largest value this sum can take on Σ, which is obtained by first assuming that every value $i \in \{1, \ldots, T\}$ appears just often enough to attain the minimum frequency $\frac{1}{T} - \epsilon$ (giving the term $\left(\frac{1}{T} - \epsilon\right) K$), and then that all additional draws ($t'(1 - (\frac{1}{T} - \epsilon)T) = t'T\epsilon$ of them) all give environment T. Comparing with (A6), (A9) is the statement that on Σ, $t'(t) < \hat{t}(t)$, and hence on Σ, all of the random variables $\zeta(1), \ldots, \zeta(t')$ are relevant.

We now put bounds on $r_\tau(t)/t$. First, note that (using (A7) for the first equality)

$$\begin{aligned} t - t' \left(\left(\frac{1}{T} - \epsilon \right) K + T\epsilon \right) &= t' \left(\left(\frac{1}{T} - \epsilon \right) K + T^2\epsilon \right) - t' \left(\left(\frac{1}{T} - \epsilon \right) K + T\epsilon \right) \\ &= t'(T^2 - T)\epsilon. \end{aligned}$$

Then, on Σ, we have

$$\frac{\rho_\tau(t')}{t} \leq \frac{r_\tau(t)}{t} \leq \frac{\rho_\tau(t') + t'(T^2 - T)\epsilon}{t}.$$

In particular, a lower bound on $r_\tau(t)$ is given by assuming that no further draws of environment τ occur past time t', giving $r_\tau(t) = r_\tau(t')$. An upper bound is given by assuming that every subsequent draw is environment τ, and that there are $t - t' \left(\left(\frac{1}{T} - \epsilon\right) K + T\epsilon \right) = t'(T^2 - T)\epsilon$ such draws.

Inserting lower and upper bounds for $\rho_\tau(t')$ (given that we are in Σ) in the appropriate places, this is (cf. (A8))

$$\frac{t' \left(\frac{1}{T} - \epsilon\right)}{t} \leq \frac{r_\tau(t)}{t} \leq \frac{t' \left(\frac{1}{T} + \epsilon\right) + (T^2 - T)\epsilon}{t}$$

and, using (A7),

$$\frac{\frac{1}{T}-\epsilon}{\left(\frac{1}{T}-\epsilon\right)K+T^2\epsilon} \le \frac{r_{\mathcal{T}}(t)}{t} \le \frac{\frac{1}{T}+\epsilon+(T^2-T)\epsilon}{\left(\frac{1}{T}-\epsilon\right)K+T^2\epsilon}.$$

There thus exist constants $0<\underline{c}<\overline{c}$ such that, for any sufficiently small ϵ and for all sufficiently large T,

$$\Pr\left\{\frac{1}{K}-\underline{c}\epsilon<\frac{r_{\mathcal{T}}(t)}{t}<\frac{1}{K}+\overline{c}\epsilon\right\}\ge 1-\epsilon$$

which implies (A5). ■

6.2 Proof of Proposition 2

The Leslie matrices identifying the two environments are:

$$\begin{array}{ll} A: & \left[\begin{array}{cc} Dx_1 & D \\ 0 & 0 \end{array}\right] \\ B: & \left[\begin{array}{cc} 0 & D \\ Dx_2 & 0 \end{array}\right]. \end{array}$$

The transition matrix between environments, M, is given by

$$\left[\begin{array}{cc} \alpha & 1-\alpha \\ 1-\alpha & \alpha \end{array}\right].$$

We then note that the stationary distribution of the matrix M attaches probability 1/2 to each environment. We consider the case in which the initial environment is drawn from this stationary distribution, so that the prior expectation for any period is also this distribution. (If the initial environment is drawn from some other distribution, we need only let the process run sufficiently long that it is almost always near the stationary distribution.) Note that

$$M^2=\left[\begin{array}{cc} \alpha^2+(1-\alpha)^2 & 2(1-\alpha)\alpha \\ 2(1-\alpha)\alpha & \alpha^2+(1-\alpha)^2 \end{array}\right]=\left[\begin{array}{cc} 1-2(1-\alpha)\alpha & 2(1-\alpha)\alpha \\ 2(1-\alpha)\alpha & 1-2(1-\alpha)\alpha \end{array}\right].$$

We now construct a backward chain. Note first

$$\begin{aligned}\Pr(s_{t-1}=A|s_t=A) &= \frac{\Pr(s_t=A|s_{t-1}=A)\Pr(s_{t-1}=A)}{\Pr(s_t=A|s_{t-1}=A)\Pr(s_{t-1}=A)+\Pr(s_t=A|s_{t-1}=B)\Pr(s_{t-1}=B)} \\ &= \frac{\alpha\frac{1}{2}}{\alpha\frac{1}{2}+(1-\alpha)\frac{1}{2}} \\ &= \alpha.\end{aligned}$$

Similarly,

$$\begin{aligned}
\Pr(s_{t-2} = A | s_t = B) &= \frac{\Pr(s_t = B | s_{t-2} = A) \Pr(s_{t-2} = A)}{\Pr(s_t = B | s_{t-2} = A) \Pr(s_{t-2} = A) + \Pr(s_t = B | s_{t-2} = B) \Pr(s_{t-2} = B)} \\
&= \frac{2(1-\alpha)\alpha\frac{1}{2}}{2(1-\alpha)\alpha\frac{1}{2} + (1 - 2(1-\alpha)\alpha)\frac{1}{2}} \\
&= 2(1-\alpha)\alpha.
\end{aligned}$$

The backward chain, giving the state in either period $t-1$ or $t-2$ as a function of the current state (the former if the current state is A, the latter if B), is then given by

$$\left[\begin{array}{cc} \alpha & 1-\alpha \\ 2(1-\alpha)\alpha & 1-2(1-\alpha)\alpha \end{array} \right].$$

We now reverse our view of the process, starting our numbering at the end, and think of this as a forward chain, giving the state in period $t+1$ as a function of the state in period t. The stationary distribution of this chain solves

$$[p, 1-p] \left[\begin{array}{cc} \alpha & 1-\alpha \\ 2(1-\alpha)\alpha & 1-2(1-\alpha)\alpha \end{array} \right] = \left[\begin{array}{c} p \\ 1-p \end{array} \right],$$

giving

$$\begin{aligned}
p\alpha + 2(1-\alpha)\alpha(1-p) &= p \\
2(1-\alpha)\alpha(1-p) &= p(1-\alpha) \\
2\alpha(1-p) &= p \\
2\alpha - 2\alpha p &= p \\
p &= \frac{2\alpha}{1+2\alpha} \\
1-p &= \frac{1}{1+2\alpha}.
\end{aligned}$$

Now we fix a time T and calculate how many draws t will be taken from the forward chain by time T, which is given by

$$\left[\frac{2\alpha}{1+2\alpha} + \frac{1}{1+2\alpha} 2 \right] t = T.$$

Our expression for the population at time T is then given by

$$\begin{aligned}
N_T &= \left(x_1^p x_2^{1-p} \right)^t \\
&= \left(x_1^{\frac{2\alpha}{1+2\alpha}} x_2^{\frac{1}{1+2\alpha}} \right) \frac{T}{\frac{2\alpha}{1+2\alpha} + \frac{2}{1+2\alpha}}
\end{aligned}$$

and hence

$$\begin{aligned}\frac{1}{T}\ln N_T &= \ln\left(x_1^{\frac{2\alpha}{1+2\alpha}}x_2^{\frac{1}{1+2\alpha}}\right)^{\frac{1+2\alpha}{2+2\alpha}}\\ &= \ln\left(x_1^{\frac{2\alpha}{2+2\alpha}}x_2^{\frac{1}{2+2\alpha}}\right)\\ &= \frac{2\alpha\ln x_1+\ln x_2}{2+2\alpha}.\end{aligned}$$

■

References

[1] George W. Ainslie. *Picoeconomics.* Cambridge University Press, Cambridge, 1992.

[2] George A. Akerlof. The market for lemons: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 89:488–500, 1970.

[3] Nabil Ibraheem Al-Najjar. Decomposition and characterization of risk with a continuum of random variables. *Econometrica*, 63(5):1195–1264, 1995.

[4] M. Allais. Le comportement de l'homme rationnnel devant le risque, critique des postulats et axiomes de l'ecole Americaine. *Econometrica*, 21:503–546, 1953.

[5] Elliott Aronson. *The Social Animal.* W. H. Freeman and Company, New York, 1995.

[6] Solomon E. Asch. Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs*, 70(416), 1956.

[7] Abhijit Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107:797–817, 1992.

[8] Abhijit Banerjee and Drew Fudenberg. Word-of-mouth learning. *Games and Economic Behavior*, 46:1–22, 2004.

[9] Jerome H. Barkow, Leda Cosmides, and John Tooby. *The Adapted Mind.* Oxford University Press, Oxford, 1992.

[10] Gary S. Becker, Kevin M. Murphy, and Ivan Werning. The equilibrium distribution of income and the market for status. *Journal of Political Economy*, 113(2):282–310, 2005.

[11] Michel Benaïm and Jörgen Weibull. Deterministic approximation of stochastic evolution in games. *Econometrica*, 71(3):873–903, 2003.

[12] Jeremy Bentham. *Principles of Morals and Legislation.* Doubleday, London, 1791.

[13] Theodore C. Bergstrom. Storage for good times and bad: Of rats and men. Mimeo, University of California, Santa Barbara, 1997.

[14] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural exchange as information cascades. *Journal of Political Economy*, 100:992–1026, 1992.

[15] Ken Binmore. Nash bargaining theory III. STICERD Discussion Paper 80/15, London School of Economics, 1980.

[16] Robert Boyd and Peter J. Richerson. *Culture and the Evolutionary Process.* University of Chicago Press, Chicago, 1985.

[17] Robert Boyd and Peter J. Richerson. Group selection among alternative evolutionarily stable strategies. *Journal of Theoretical Biology*, 145:331–342, 1990.

[18] Philip Brickman, Dan Coates, and Ronnie Janoff-Bulman. Lottery winners and accident victims: Is happiness relative? *Journal of Personality and Social Psychology*, 32(8):917–927, 1978.

[19] Michael Bulmer. *Theoretical Evolutionary Ecology.* Sinauer Associates, Inc., Sunderland, MA, 1997.

[20] Terry Burnham and Jay Phelan. *Mean Genes.* Perseus Publishing, Cambridge, 2000.

[21] Colin Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction.* Russell Sage Foundation and Princeton University Press, Princeton, 2003.

[22] Colin Camerer, George Loewenstein, and Matthew Rabin, editors. *Advances in Behavioral Economics.* Russell Sage Foundation, New York, 2004.

[23] Brian Charlesworth. *Evolution in Age-Structured Populations.* Cambridge University Press, Cambridge, 1994.

[24] Robert B. Cialdini. *Influence: Science and Practice.* Scott, Foresman and Company, Boston, 1988.

[25] Andrew Clark. *Microcognition.* MIT Press, Cambridge, Massachusetts, 1993.

[26] Harold L. Cole, George J. Mailath, and Andrew Postlewaite. Social norms, savings behavior, and growth. *Journal of Political Economy*, 100:1092–1125, 1992.

[27] William S. Cooper and Robert H. Kaplan. Adaptive "coin-flipping": A decision theoretic examination of natural selection for random individual variation. *Journal of Theoretical Biology*, 94(1):135–151, 2004.

[28] James C. Cox and Vjollca Sadiraj. Small- and large-stakes risk aversion: Implications of concavity calibration for decision theory. *Games and Economic Behavior*, 56(1):45–60, 2006.

[29] Vincent P. Crawford and Hal Varian. Distortion of preferences and the Nash theory of bargaining. *Economics Letters*, 3:203–206, 1979.

[30] Philip A. Curry. Decision making under uncertainty and the evolution of interdependent preferences. *Journal of Economic Theory*, 98(2):357–369, 2001.

[31] Partha Dasgupta and Eric Maskin. Uncertainty and hyperbolic discounting. *American Economic Review*, 95(4):1290–1299, 2005.

[32] Richard Dawkins. *The Selfish Gene.* Oxford University Press, Oxford, 1989.

[33] Anna Dreber and Moshe Hoffman. Portfolio selection in utero. Mimeo, Stockholm School of Economics and University of Chicago, 2007.

[34] J. S. Duesenberry. *Income, Saving, and the Theory of Consumer Behavior.* Harvard University Press, Cambridge, Massachussetts, 1949.

[35] D. Einon. How many children can one man have? *Evolution and Human Behavior*, 19:413–426, 1998.

[36] Glenn Ellison and Drew Fudenberg. Rules of thumb for social learning. *Journal of Political Economy*, 101:612–643, 1993.

[37] Glenn Ellison and Drew Fudenberg. Word-of-mouth communication and social learning. *Quarterly Journal of Economics*, 110:95–126, 1995.

[38] Daniel Ellsberg. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75:643–669, 1961.

[39] Jon Elster. Weakness of will and the free-rider problem. *Economics and Philosophy*, 1:231–265, 1985.

[40] Jeffrey C. Ely and Okan Yilankaya. Nash equilibrium and the evolution of preferences. *Journal of Economic Theory*, (2):255–272, 2000.

[41] I. Eshel. On the neighbor effect and the evolution of altruistic traits. *Theoretical Population Biology*, 3:258–277, 1972.

[42] Jesus Fernandez-Villaverde and Arijit Mukherji. Can we really observe hyperbolic discounting? Mimeo, University of Pennsylvania, 2001.

[43] Irving Fisher. *The Theory of Interest, as Determined by Impatience to Spend Income and Opportunity to Invest It.* MacMillan, New York, 1930.

[44] Hugh J. Foley and Margaret W. Matlin. *Sensation and Perception.* Allyn and Bacon, Boston, 2009.

[45] R. H. Frank. If homo economicus could choose his own utility function, would he want one with a conscience? *American Economic Review*, 77:593–604, 1987.

[46] Robert H. Frank. *Passions within Reason.* Norton, New York, 1988.

[47] Robert H. Frank. *Luxury Fever.* Free Press, New York, 1999.

[48] Shane Frederick and George Loewenstein. Hedonic adaptation. In Edward Diener Danield Kahneman and Norbert Schwartz, editors, *Well-Being: The Foundations of Hedonic Psychology*, pages 302–329. Russell Sage Foundation Press, New York, 1999.

[49] Shane Frederick, George Loewenstein, and Ted O'Donoghue. Time discounting and time preference: A critical view. *Journal of Economic Literature*, 40(2):351–401, 2002.

[50] Bruno S. Frey and Alois Stutzer. *Happiness and Economics: How the Economy and Institutions Affect Human Well-Being.* Princeton University Press, Princeton, 2002.

[51] Bruno S. Frey and Alois Stutzer. What can economists learn from happiness research? *Journal of Economic Literature*, 40:402–435, 2002.

[52] Daniel Friedman. The S-shaped value function as a contrained optimum. *American Economic Review*, 79(5):1243–1248, 1989.

[53] Milton Friedman. Choice, chance, and the personal distribution if income. *Journal of Political Economy*, 61(4):277–290, 1953.

[54] Milton Friedman and Leonard J. Savage. The utility analysis of choices involving risk. *Journal of Political Economy*, 56(4):279–304, 1948.

[55] Drew Fudenberg and David K. Levine. *Theory of Learning in Games.* MIT Press, Cambridge, 1998.

[56] H. Furstenberg and H. Kesten. Products of random matrices. *Annals of Mathematical Statistics*, 31(2):457–469, 1960.

[57] E. L. Gardner and J. H. Lowinson. Drug craving and positive/negative hedonic brain substrates activated by addicting drugs. *Seminars in the Neurosciences*, 5:359–368, 1993.

[58] Gerd Gigerenzer. *Calculated Risks.* Simon and Schuster, New York, 2002.

[59] Daniel Gilbert. *Stumbling on Happiness.* Vintage Books, New York, 2007.

[60] Itzhak Gilboa and Larry Samuelson. Subjectivity in inductive inference. Cowles Foundation Discussion Paper 1725, Tel Aviv University and Yale University, 2009.

[61] S. J. Gould and R. C. Lewontin. The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptionist programme. *Proceedings of the Royal Society of London, Series B*, 205:581–598, 1979.

[62] Alan Grafen. Modelling in behavioural ecology. In J. R. Krebs and N. B. Davies, editors, *Behavioral Ecology: An Evolutionary Apoproach*, pages 5–31. Blackwell Scientific Publications, Oxford, 1991.

[63] Alan Grafen. Formal Darwinism, the individual-as-maximizing-agent analogy and bet-hedging. *Proceedings of the Royal Society of London Series B*, 266:799–803, 1999.

[64] Faruk Gul and Wolfgang Pesendorfer. Temptation and self-control. *Econometrica*, 69(6):1403–1436, 2001.

[65] Faruk Gul and Wolfgang Pesendorfer. Mindless economics. In Andrew Caplin and Andrew Shotter, editors, *The Foundations of Positive and Normative Economics*. Oxford University Press, New York, 2008.

[66] Werner Güth. An evolutionary approach to explaining cooperative behavior by reciprocal incentives. *International Journal of Game Theory*, 24:323–344, 1995.

[67] Werner Güth and Menahem E. Yaari. Explaining reciprocal behavior in simple strategic games: An evolutionary approach. In Ulrich Witt, editor, *Explaining Process and Change*, pages 23–34. University of Michigan Press, Ann Arbor, 1992.

[68] Ingemar Hansson and Charles Stuart. Malthusian selection of preferences. *American Economic Review*, 80(3):529–544, 1990.

[69] David G. C. Harper. Communication. In J. R. Krebs and N. B. Davies, editors, *Behavioural Ecology*, pages 374–397. Blackwell Scientific Publications, London, 1991.

[70] P. H. Harvey, R. D. Martin, and T. H. Clutton-Brock. Life histories in comparative perspective. In Barbara B. Smuts, Dorothy L. Cheney, Robert M. Seyfarth, Richard W. Wrangham, and Thomas T. Struhsaker, editors, *Primate Societies*, pages 181–196. University of Chicago Press, Chicago, 1986.

[71] Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, and Herbert Gintis. *Foundations of Human Sociality.* Oxford University Press, Oxford, 2004.

[72] Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, and Richard McElreath. In search of Homo Economicus: Behavioral experiments in 15 small-scale societies. *American Economic Review*, 91 (2):73–78, 2001.

[73] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics.* Cambridge University Press, Cambridge, 1998.

[74] Ed Hopkins and Tatiana Kornienko. Running to keep in the same place: Consumer choice as a game of status. *American Economic Review*, 94(4):1085–1107, 2004.

[75] Alasdair I. Houston and John M. McNamara. *Models of Adaptive Behavior*. Cambridge University Press, Cambridge, 1999.

[76] Alex Kacelnik. Normative and descriptive models of decision making: Time discounting and risk sensitivity. In Gregory R. Bock and Gail Cardew, editors, *Characterizing Human Psychological Adaptations*, pages 51–70. John Wiley and Sons, New York, 1997.

[77] Daniel Kahneman and Amos Tversky. Judgement under uncertainty: Heuristics and biases. In Daniel Kahneman, Paul Slovic, and Amos Tversky, editors, *On the Psychology of Prediction*, pages 48–68. Cambridge University Press, Cambridge, 1982.

[78] A. Knafo, S. Israel, A. Darvasi, R. Bachner-Melman, F. Uzefovsky, L. Cohen, E. Feldman, E. Lerer, E. Laiba, Y. Raz, L. Nemanov, I. Gritsenko, C. Dina, G. Agam, B. Dean, G. Bronstein, and R. P. Ebstein. Individual differences in allocation of funds in the dictator game associated with length of the arginine vasopressin 1a receptor RS3 promoter region and correlation between RS3 length and hippocampal mRNA. *Genes, Brain and Behavior*, 2007. OnlineEarly Article, doi:10.1111/j.1601-183X.2007.00341.x.

[79] David Laibson. A cue-theory of consumption. *Quarterly Journal of Economics*, 116(1):81–120, 2001.

[80] Joseph LeDoux. *The Emotional Brain*. Simon and Schuster, New York, 1996.

[81] P. H. Leslie. On the use of matrices in certain population mathematics. *Biometrica*, 33(3):183–212, 1945.

[82] P. H. Leslie. Some further notes on the use of matrices in population mathematics. *Biometrica*, 35(1–2):213–245, 1948.

[83] W. Leutenegger. Encephalization and obstetrics in primates with particular reference to human evolution. In E. Armstrong and D. Falk, editors, *Primate Brain Evolution: Methods and Concepts*, pages 85–95. Plenum Press, New York, 1982.

[84] R. G. Lipsey and Kelvin Lancaster. The general theory of second best. *Review of Ecohnomic Studies*, 24(1):11–32, 1956.

[85] George Loewenstein. Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes*, 65:272–292, 1996.

[86] George Loewenstein and Drazen Prelec. Anomalies in intertemporal choice: Evidence and an interpretation. *Quarterly Journal of Economics*, 107:573–598, 1992.

[87] George Loewenstein and David Schkade. Wouldn't it be nice? Predicting future feelings. In Daniel Kahneman, Ed Diener, and Norbert Schwarz, editors, *Well-Being: The Foundations of Hedonic Psychology*, pages 85–105. Russell Sage Foundation, New York, 1999.

[88] George Loewenstein and Richard Thaler. Anomalies: Intertemporal choice. *Journal of Economic Perspectives*, 3:181–193, 1989.

[89] George J. Mailath. Do people play Nash equilibrium? Lessons from evolutionary game theory. *Journal of Economic Literature*, 36:1347–1374, 1998.

[90] John Maynard Smith. Group selection and kin selection. *Nature*, 201:1145–1147, 1964.

[91] John Maynard Smith. *Evolutionary Genetics, 2nd Edition.* Oxford University Press, Oxford, 1998.

[92] J. G. Mercer, C. L. Adam, and P. J. Morgan. Towards an understanding of physiological body mass regulation: Seasonal animal models. *Nutritional Neuroscience*, 3:307–320, 2000.

[93] K. Milton. Foraging behavior and the evolution of primate cognition. In R. W. Byrne and A. Whiten, editors, *Machiavellian Intellegence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans*, pages 285–305. Clarendon Press, Oxford, 1988.

[94] S. Mineka and M. Cook. Mechanisms involved in the operational conditioning of fear. *Journal of Experimental Psychology: General*, 122:23–38, 1993.

[95] W. Mischel, Y. Shoda, and M. L. Rodriguez. Delay of gratification in children. In George Loewenstein and Jon Elster, editors, *Choice over Time*. Russell Sage, New York, 1992.

[96] Robert F. Nau and Kevin F. McCardle. Coherent behavior in noncooperative games. *Journal of Economic Theory*, 50(2):424–444, 1990.

[97] Nick Netzer. Evolution of time preferences and attitudes towards risk. *American Economic Review*, 99(3):937–55, 2009.

[98] David Neumark and Andrew Postlewaite. Relative income concerns and the rise in married women's employment. *Journal of Public Economics*, 70:157–183, 1998.

[99] Georg Nöldeke and Larry Samuelson. Information-based relative consumption effects: Correction. *Econometrica*, 73:1383–1387, 2005.

[100] Ted O'Donoghue and Matthew Rabin. Doing it now or later. *American Economic Review*, 89(1):103–124, 1999.

[101] Ted O'Donoghue and Matthew Rabin. Incentives for procrastinators. *Quarterly Journal of Economics*, 114(3):769–816, 1999.

[102] Efe Ok and Fernando Vega-Redondo. On the evolution of individualistic preferences: An incomplete information scenario. *Journal of Economic Theory*, 97(2):231–254, 2000.

[103] Elinor Ostrom. Collective action and the evolution of social norms. *Journal of Economic Perspectives*, 14 (3):137–158, 2000.

[104] A. T. Peacock. *The Theory of the Market Economy*. William Hodge, London, 1952.

[105] Clive J. Petry and C. Nicholas Hales. Long-term effects on offspring of intrauterine exposure to deficits in nutrition. *Human Reproduction Update*, 6:578–586, 2000.

[106] Steven Pinker. *How the Mind Works*. W. W. Norton, New York, 1997.

[107] R. Pluchik. Emotions: A general psychoevolutionary theory. In K. R. Scherer and P. Ekman, editors, *Approaches to Emotion*, pages 197–219. Erlbaum, Hillsdale, NJ, 1984.

[108] Robert A. Pollak. Consistent planning. *Review of Economic Studies*, 35(2):201–208, 1968.

[109] Andrew Postlewaite. The social basis of interdependent preferences. *European Economic Review*, 42(3–5):779–800, 1998.

[110] Matthew Rabin. Psychology and economics. *Journal of Economic Literature*, 36:11–46, 1998.

[111] Matthew Rabin. Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68:1281–1292, 2000.

[112] Anita C. J. Ravelli, Jan H. P. van der Meulen, Clive Osmond, David J. P. Barker, and Otto P. Bleker. Obesity at the age of 50 y in men and women exposed to famine prenatally. *American Journal of Clinical Nutrition*, 70:811–816, 1999.

[113] Debraj Ray and Arthur J. Robson. Status, intertemporal choice, and risk-taking. Technical report, New York University and Simon Fraser University, 2010.

[114] Luis Rayo and Gary Becker. Evolutionary efficiency and happiness. *Journal of Political Economy*, 115(2):302–337, 2007.

[115] Matt Ridley. *The Red Queen.* Penguin Books, New York, 1993.

[116] Matt Ridley. *Nature via Nurture: Genes, Experience, and What Makes Us Human.* Harper Collins Publishers, New York, 2003.

[117] Arthur J. Robson. Efficiency in evolutionary games: Darwin, Nash, and the secret handshake. *Journal of Theoretical Biology*, 144:379–396, 1990.

[118] Arthur J. Robson. Status, the distribution of wealth, private and social attitudes to risk. *Econometrica*, 60(4):837–857, 1992.

[119] Arthur J. Robson. A biological basis for expected and non-expected utility. *Journal of Economic Theory*, 68(2):397–424, 1996.

[120] Arthur J. Robson. The evolution of attitudes to risk: Lottery tickets and relative wealth. *Games and Economic Behavior*, 14:190–207, 1996.

[121] Arthur J. Robson. The biological basis of economic behavior. *Journal of Economic Literature*, 39(1):11–33, 2001.

[122] Arthur J. Robson. Why would nature give individuals utility functions? *Journal of Political Economy*, 109(4):900–914, 2001.

[123] Arthur J. Robson. Group selection. In Steven N. Durlauf and Lawrence E. Blume, editors, *New Palgrave Dictionary of Economics.* Palgrave Macmillan, New York, 2008.

[124] Arthur J. Robson and Larry Samuelson. The evolution of intertemporal preferences. *American Economic Review*, 97(2 (May)):496–500, 2007.

[125] Arthur J. Robson and Larry Samuelson. The evolution of time preference with aggregate uncertainty. *American Economic Review*, 99:1925—1953, 2009.

[126] Arthur J. Robson and Larry Samuelson. The evolutionary optimality of decision and experienced utility. Working paper, Simon Fraser University and Yale University, 2010.

[127] Arthur J. Robson and Balazs Szentes. Evolution of time preference by natural selection: Comment. *American Economic Review*, 98(3):1178–1188, 2008.

[128] Arthur J. Robson, Balazs Szentes, and Emil Iantchev. The evolutionary basis of time preference: Intergenerational transfers and sex. Mimeo, Simon Fraser University and the London School of Economics, 2010.

[129] Alan R. Rogers. Evolution of time preference by natural selection. *American Economic Review*, 84(2):460–481, 1994.

[130] Ariel Rubinstein. "Economics and Psychology"?: The case of hyperbolic discounting. *International Economic Review*, 44:1207–1216, 2003.

[131] Larry Samuelson. *Evolutionary Games and Equilibrium Selection.* MIT Press, Cambridge, 1997.

[132] Larry Samuelson. Introduction to the evolution of preferences. *Journal of Economic Theory*, 97:225–230, 2001.

[133] Larry Samuelson. Information-based relative consumption effects. *Econometrica*, 72(1):93–118, 2004.

[134] Larry Samuelson and Jeroen Swinkels. Information and the evolution of the utility function. *Theoretical Economics*, 1:119–142, 2006.

[135] Paul Samuelson. A note on the measurement of utility. *Review of Economic Studies*, 4:155–161, 1937.

[136] Leonard J. Savage. *The Foundations of Statistics.* Dover Publications, New York, 1972 (originally 1954).

[137] Thomas Schelling. *The Strategy of Conflict*. Harvard University Press, Cambridge, MA., 1980 (first edition 1960).

[138] Thomas Schelling. Self-command in practice, in policy, and in a theory of rational choice. *American Economic Review*, 74(1):1–11, 1984.

[139] Paul Seabright. *The Company of Strangers: A Natural History of Economic Life*. Princeton University Press, Princeton, 2005.

[140] E. Seneta. *Non-Negative Matrices and Markov Chains*. Springer Verlag, New York, 1981.

[141] Eldar Shafir, Itamar Simonson, and Amos Tversky. Reason-based choice. *Cognition*, 49:11–36, 1993.

[142] S. Siegel. The role of conditioning in drug tolerance and addiction. In J. D. Keehn, editor, *Psychopathalogy in Animals: Research and Treatment Implications*. Academic Press, New York, 1979.

[143] Paul Slovic. *The Perception of Risk*. Earthscan Publications, London, 2000.

[144] Paul Slovic, Baruch Fishhoff, and Sarah Lichtenstein. Why study risk perception? *Risk Analysis*, 2(2):83–93, 1982.

[145] Trenton G. Smith. The McDonald's equilibrium: Advertising, empty calories, and the endogenous determination of dietary preferences. *Social Choice and Welfare*, 23(3):383–413, 2004.

[146] Elliott Sober and David Sloan Wilson. *Unto Others*. Harvard University Press, Cambridge, Massachusetts, 1998.

[147] Peter D. Sozou. On hyperbolic discounting and uncertain hazard rates. *Proceedings of the Royal Society of London, Series B*, 265(1409):2015–2020, 1998.

[148] Stephen C. Stearns and Rolf F. Hoekstra. *Evolution: An Introduction*. Oxford University Press, Oxford, second edition, 2005.

[149] Robert H. Strotz. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23(3):165–180, 1956.

[150] David Tanny. On multitype branching processes in a random environment. *Advances in Applied Probability*, 13(3):464–497, 1981.

[151] Richard H. Thaler. *Quasi-Rational Economics.* Russell Sage Foundation, New York, 1994.

[152] Richard H. Thaler and H. M. Shefrin. An economic theory of self-control. *Journal of Political Economy*, 89(2):392–406, 1981.

[153] Leon Trémblay and Wolfram Schultz. Relative reward preference in primate orbitofrontal cortex. *Nature*, 398:704–708, 1999.

[154] Amos Tversky and Itamar Simonson. Context-dependent preferences. *Management Science*, 39:1179–1189, 1993.

[155] Eric van Damme. *Stability and Perfection of Nash Equilibria.* Springer-Verlag, Berlin, 1991.

[156] Thorstein Veblen. *The Theory of the Leisure Class.* MacMillan, New York, 1899.

[157] Fernando Vega-Redondo. *Evolution, Games, and Economic Behavior.* Oxford University Press, Oxford, 1996.

[158] Heinrich von Stackelberg. *Marktform und Gleichgewicht.* Julius Springer, Vienna, 1934.

[159] Jörgen W. Weibull. *Evolutionary Game Theory.* MIT Press, Cambridge, 1995.

[160] G. C. Williams. *Adaptation and Natural Selection.* Princeton University Press, Princeton, 1966.

[161] Margo Wilson and Martin Daly. Life expectancy, economic inequality, homicide, and reproductive timing in Chicago neighborhoods. *British Medical Journal*, 314(7089):1271, 1997.

[162] Robert B. Wilson. Reputations in games and markets. In Alvin E. Roth, editor, *Game-Theoretic Models of Bargaining*, pages 27–62. Cambridge University Press, Cambridge, 1985.

[163] David H. Wolpert and David S. Leslie. The effects of observational limitations on optimal decision making. NASA Ames Research Center and Department of Mathematics, University of Bristol, 2009.

[164] V. C. Wynne-Edwards. *Animal Dispersion in Relation to Social Behavior.* Oliver and Boyd, Edinburgh, 1962.

[165] V. C. Wynne-Edwards. *Evolution Through Group Selection.* Blackwell Scientific Publications, Oxford, 1986.

[166] Menahem E. Yaari. Uncertain lifetime, life insurance, and the theory of the consumer. *Review of Economic Studies*, 32(1):137–158, 1965.

[167] Peyton Young. *Individual Strategy and Social Structure.* Princeton University Press, Princeton, 1998.

[168] Kareem A. Zaghloul, Justin A. Blanco, Christoph T. Weidemann, Kathryn McGill, Jurg L. Jaggi, Gordon H. Baltuch, and Michael J. Kahana. Human substantia nigra neurons encode unexpected financial rewards. *Science*, 323:1496–1499, 2009.