BIOLOGICAL FOUNDATIONS OF ECONOMIC BEHAVIOR[1]

Arthur J. Robson[2] ⓡ Nikolaus Robalino [3]

June 2025

ABSTRACT

We advocate here consideration of the biological origins of economic behavior. As a first approach, we look at models of adaptive and possibly hedonic utility that are inspired by neuroscience. We sketch evidence concerning how utility is implemented in the human brain and its deep evolutionary pedigree, going back beyond hunting and gathering. We show that these models make predictions relevant to modern behavior. As a second approach, we then turn to essentially demographic models for the evolution of time preference and for attitudes to risk. Hunting and gathering was the crucible in which these key aspects of preferences were refined. We show that evolution might have produced empirically plausible pure rates of time preference in hunter-gatherer societies, despite the lack of opportunities to save or borrow. Also as part of the demographic approach, we show that there might be greater aversion to aggregate risk than to idiosyncratic risk and look at modern experimental evidence confirming this. We consider theory of mind as a factor favoring the operation of hunter-gatherer societies and discuss a demographic model of a key aspect of this—the ability to ascribe utility to others in strategic settings. We consider modern experimental evidence connecting such behavior with autism, taken as the lack of theory of mind. JEL Codes: A12, D11.

1. INTRODUCTION

The traditional economic approach to preferences is "creationist". That is, a set of abstract axioms is proposed, evaluated in terms of their inherent plausibility, and

---

[1] We thank Ben Enke, Nathan Nunn, Larry Samuelson, and participants at the conference at Northwestern University on the Economics of Culture for helpful comments. We thank Babak Sahragard for research assistance. Robson thanks the Social Sciences and Humanities Research Council of Canada for financial support. The authors names were randomized by the AEA author randomization tool https://www.aeaweb.org/journals/random-author-order.

[2] Department of Economics, Simon Fraser University, Burnaby BC Canada V5A 1S6. robson@sfu.ca.

[3] Department of Economics, Rochester Institute of Technology, Rochester, NY USA 14623. ndrobalin@gmail.com.

the logical implications for choice behavior of the axioms are derived.[4]An evolutionary approach considers instead how particular preferences might realistically have been the product of Darwinian natural selection, so outcompeting and displacing the alternatives.

The evolutionary approach does not inevitably contradict the axiomatic approach but suggests modifications that might not otherwise suggest themselves. For example, an essentially demographic theory of the evolution of attitudes to risk implies classical expected utility when gambles are statistically independent across individuals. That is, individuals whose choices maximize the appropriate form of expected utility will numerically swamp those whose choices do not. When gambles are correlated, however, the favored risk attitudes may no longer be consistent with expected utility, but might instead entail interdependent preferences. This latter possibility arises specifically due to the demographic nature of the model and might seem implausible outside of an evolutionary context.

The evolutionary approach takes little for granted. It is not merely the properties of utility that are subjected to the forge of natural selection, but aspects of decision making that are often taken for granted. Indeed, Section 2 begins by asking the basic question—Why might we have utility functions in the first place? The answer given is that the possession of a utility function allows the individual free agency, so individuals can take full advantage of local information.

Given a basic rationale for utility functions, we then move on to ask what properties they might have. Robson and Samuelson (2024), for example, argue that neuroscience provides evidence that utility is not merely a pale reflection of choice as in the theory of revealed preference but is “substantive”—a term defined by them to mean represented as quantifiable neural activity in the brain, where this activity drives choice rather than simply reflecting it.

---

[4] For example, von Neumann and Morgenstern (1944) derive expected utility theory from a small set of reasonable axioms on choice in the presence of uncertainty. Savage (1954) extends this work to account for subjective probabilities.

This neural perspective is the first of two evolutionary approaches considered here.[5] Neuroscientiic evidence first implicates hard-wired cognitive structures as underlying decision-making. These then suggest features to be incorporated in models of utility that are ultimately biological in origin. This approach is not maximally reductionist, but is built on a solid foundation of direct observation.

Since all neural processes are irreducibly noisy, this implies that utility should adapt to circumstances. In this chapter, we review a model of this adaptation and its empirical implications—a preference for growth, the "hedonic treadmill" and the Easterlin paradox, among these. The rich implications constitute the "proof of the pudding" for the adaptive model.

The resulting model of adaptive substantive utility indeed has some kinship to the large literature on habit formation in economics. For example, Campbell and Cochrane (1999) apply a simple model of habit formation to address anomalies arising in dynamic asset pricing. The agent evaluates present consumption relative to a stock representing the level of habit, where this level is itself built up from the history of consumption. However, the details of the present formulation are new; further, the dependence on history here is a result rather than an assumption.

The second approach to the evolutionary basis of behavior examined in this chapter is explicitly evolutionary and essentially demographic in nature. This considers the effect of various choices on population growth rates, using demographic models of the implications for growth rates of lifetime profiles of fertility and mortality. The type having the lifetime profiles that induce the fastest growth will then be selected since it will ultimately dominate a population that starts with any mixture of the types. This represents a Darwinian methodology in a straightforward manner. A systematic, explicitly evolutionary approach like this to preferences seems largely new in economics.

Section 3 investigates the implications of this demographic approach for attitudes to idiosyncratic risk and time preference. It first shows that an individual should

---

[5] The second approach taken in this chapter was informally applied already in this introduction. More formally, it relies on demographic modeling and an explicit analysis of evolutionary dynamics. We discuss this second approach in greater detail in Sections 3 and 4.

be risk neutral with respect to idiosyncratic risk concerning offspring, which translates readily into the expected utility theorem. We argue for the plasticity of attitudes to risk so that the choice of gamble might adjust appropriately to the environment, given this can be observed by the individual. Such plasticity arises as the careful tailoring of utility to current circumstances rather than a capricious change in attitudes to risk.

This demographic approach generates the simplest standard form of time preference, where an intertemporal consumption stream is evaluated as a sum of flow utilities weighted by a geometrically growing discount factor. A geometrically growing contribution to the discount factor arises from a new source—population growth. There is also a contribution from mortality to discounting, which is a well-known effect, being described by Fisher (1930).

Such time preference might have been exercised in hunter-gatherer societies despite the lack of opportunities to borrow or save, involving instead the substantial intergenerational transfers then prevalent. The pure rate of time preference found theoretically is shown to be capable of accounting for moderate modern observations. This rate of time preference may also be plastic—adjusting appropriately to variation in the observed environment. Furthermore, the model suggests testable hypotheses concerning how the ancestral environment of a particular country would influence modern rates of time preference.

Section 4 then shows how this second demographic approach generates a greater aversion to aggregate risk than to idiosyncratic risk. In the model of Robson (1996), individuals who, in each generation, experience a distribution of offspring independently of others do better in the long run than do individuals who experience the same distribution but where the outcomes are correlated across individuals. This greater aversion to aggregate risk than to idiosyncratic risk is not an hypothesis that seems to have arisen previously in just that guise.

Robson (1996) has generated a small but persistent literature. The role of the discrete time basis of the effect was illuminated by Robatto and Szentes (2013) who show that, in a continuous time framework where the the birth rate and mortality are age independent, the contrast between attitudes to aggregate and idiosyncratic risk disappears. Robson and Samuelson (2019) show that the contrast generally reappears in a continuous time framework in which fertility and mortality are age dependent, although it is not inevitable that idiosyncratic risk is preferred. Robson and Samuelson (2024) present an extensive discussion of this theoretical debate.

Here, we consider theoretical arguments that are helpful in framing the contrast empirically—and review work by Oprea and Robalino (2024) showing that such a distinction can be seen in experiments. These experiments show that a significant distinction arises even in a laboratory setting where aggregate risk entails outcomes that are shared only by the subjects actually present.

This surprising prediction concerning a differential response to aggregate risk drawn from an evolutionary perspective dramatizes the scientific value of this approach.

The focus so far has been on the individual. Ultimately, however, it is crucial to consider human social interactions, which are often the central concern. We take a step in this direction in Section 5 by looking at "theory of mind" (ToM) as in psychology where agents recognize the agency of themselves and others. Individuals have only a "theory" of another individual's preferences and beliefs, for example, since they must infer these from behavior. This is perhaps the central component of human social intelligence. Here we consider, in particular, the implications of recognizing that others have preferences that may differ from one's own.

Using Robalino and Robson (2016) and Kimbrough, Robalino and Robson (2017), we discuss theoretically why the attribution of payoffs to others would promote evolutionary success. This applies essentially the demographic approach. We then look at relevant experiments from Kimbrough, Robalino ,and Robson (2017) that examine how well individuals predict the choices of others in strategic interactions. If you have observed another agent's choice between various alternatives in previous plays of a game, you may be able to predict their choice in future plays. The game is constructed so that this prediction can be made with varying degrees of sophistication. We correlate performance in such strategic settings with measures of autism, a key component of which is thought to be a lack of theory of mind.

2. A BIOLOGICAL RATIONALE FOR UTILITY

We start with a broad evolutionary rationale for a fundamental economic concept—that of utility—a rationale that might apply to distant ancestral species of humans as well. The basic framework we propose is that of the principal-agent problem, where evolution is the principal and the individual is the agent.[6]

A typical application in economics for the principal-agent model is to the interaction between an owner of a firm and the manager. The manager knows much more about the day to day operation of the firm, knows how hard he himself works, in particular. The owner observes the profits of the firm which is only a noisy signal of the manager's effort, and designs a contract to best incentivize the manager to maximize the net profit of the owner.

As in that model, there is a gain from evolution devolving partial control to the individual because the individual has information about current local conditions that evolution does not. More precisely, this is because current local conditions are so novel that evolution would not have been able to program an optimal response, to build an automaton that responds appropriately to these local conditions.

Why would it make sense for evolution to devolve only some control over the actions of the individual? The wrinkle here is that evolution also knows things that the individual does not. These might concern the long-run evolutionary consequences of outcomes attained. This consideration means that it might well be optimal for evolution to convey such information by means of a utility function. For example, the fitness advantages of a balanced diet might not be immediately apparent to an individual, who only has limited direct data available. However, if the appropriate utility function were embedded in the agent, the experience of many generations of ancestors concerning a balanced diet could be effectively taken into account.

Consider a simple model to illustrate this argument concerning dietary diversity. Suppose an individual must first choose one of I foods to forage for, and then one

---

[6] This analogy to the principal-agent model was proposed by Binmore (1994). The principal agent model in economics is outlined in Mas-Colell, Whinston and Green (1995, Chapter 14).

of J locations at which to forage. Choosing good $i = 1, ..., I$ and location $j = 1, ..., J$ means the individual has probability $p^i_j \in [0, 1]$ of obtaining 1 unit of this i-th good in one period.[7] We will consider two cases—either that these probabilities are known by evolution or—if they are not—that they are instead known to the individual.

To simplify the calculations, suppose draws for these goods at these locations are taken frequently, so that there are r draws per unit of time, with the quantities of the various goods scaled accordingly by $1/r$ to keep the overall level of consumption comparable. As $r \to \infty$, the choice of good i at location j then generates a deterministic flow of that good at rate p^i_j, by the strong law of large numbers.

Dietary diversity is introduced by supposing the individual can apportion her time across the possible goods and locations. Suppose the proportion of time the individual chooses good i is π_i for each i, where $\sum\limits_i \pi_i = 1$ and suppose that the location used if good i is chosen is j_i.[8] Suppose the arrival rate of offspring is then given by the function

$$\lambda \left(\pi_1 p^1_{j_1}, \pi_2 p^2_{j_2}, ... \right)$$

which reflects how flow fitness depends on the flow rates of goods. This function λ is assumed to be strictly increasing in each argument.

Let us first suppose that the various probabilities p^i_j are known to evolution, as is the utility function λ, meaning more precisely just that this entire problem has persisted in its current form for a long time. Then evolution's design problem is straightforward—solve

$$\max_\pi \lambda \left(\pi_1 \max_j p^1_j, \pi_2 \max_j p^2_j, ... \right),$$

since choosing j_i to maximize p^i_j for each i is uniformly best, regardless of π. That is, in this case, evolution can reduce the role of the individual, making the evolutionarily appropriate choices on his behalf of time allocation π_i and of location conditional on the good chosen. The assumption that evolution knows all these

---

[7] Obtaining a good i at date t is independent of obtaining other goods $j \neq i$ and over t.

[8] Mixing over locations would not be difficult to allow.

probabilities, or that there has been sufficient time for the appropriate optimal choices to have been selected, is not generally plausible, however.

Suppose now, more plausibly, that evolution does not know the probabilities p_j^i. The entire problem is too transient for evolution to exploit. In particular, there has not been time for the optimal choice of good and location to have been directly prescribed by evolution. The probabilities can, however, be observed by the individual, who is familiar with local idiosyncrasies.

This situation creates an opportunity for a mechanism of partial devolution of control to be evolutionarily optimal. Since evolution does not know the probabilities, it is desirable to allow the individual to make use of her knowledge of these. One way to achieve the optimum is to embed evolution's function λ—flow fitness—as the utility function in each agent. In addition, agents should be given the motivation to maximize this utility function. The agent is then responsible for picking the best location for each good and for choosing the right weights for each good given the utility.

A conceptually related model is in Robson (2001) . In that model the individual faces a two armed bandit that becomes degenerate with many repetitions. It then pays for evolution to dictate the von Neumann-Morgenstern utility of the various outcomes, while leaving the individual to learn the appropriate probabilities.

Why does the optimal decentralization of control from evolution to the individual stop here? Why give the individual a utility function that depends on consumption? Why wouldn't the appropriate utility be simply offspring? Why not have the individual simply observe the actual offspring produced as a consequence of any choice of consumption? Offspring, after all, are only the final good, and consumption is just a means to an end, an intermediate good. Why do we care about food, sex, and so on, when these are only biologically advantageous for the offspring they help produce?

One answer is as follows. The problem is that there is only a small number of offspring produced in each lifetime. If a population were in steady state, each parthenogenetic (asexual) individual would produce one offspring on average. With sex, each couple would produce two, in the steady state. Not many, in any event. If an individual relied on observing the actual number of offspring produced, there would necessarily be a significant probability of making the wrong choice. Relying on the right utility would reduce the probability of such mistakes to zero.

Suppose, for example, that an individual faces two options, to strip the problem to the bare minimum. Option A has higher expected offspring than option B and hence should be chosen. In particular, option A generates 1 offspring with probability $1/3$ and 2 offspring with probability $2/3$ and option B generates 1 offspring with probability $2/3$ and 2 with probability $1/3$.[9] The individual does not know this and has no reason to prefer either option over the other, but can observe one independent draw of each option and then has to commit to one of them. If one option generates more offspring in the trial, that option is chosen, if they generate the same number, choice is made on a 50-50 basis. The probability of committing to A is then $(2/3)(2/3) + 2(1/2)(1/3)(2/3) = 2/3$ which is the probability of observing 2 offspring from A and 1 from B plus a half the probability of observing the same number from A and B. Hence the probability of committing to B is $1/3 > 0$ which is less than the probability of committing to A, but still positive.

Building a preference for A into the individual's utility, on the other hand, can be accomplished by using the evolutionary experience of a large number of ancestors and ensures the correct choice can be always made. More precisely, the individuals who give a higher utility to A than to B have expected offspring $1/3 + 4/3 = 5/3$ and hence grow faster than the individuals who give a higher utility to B than to A and have expected offspring $2/3 + 2/3 = 4/3$. Clearly, the type that prefers A does better than the type that experiments, since B is chosen by $1/3$ of these types. In this sense, evolution is wise to things that the individual is not.

2.1. **The Neural Representation of Utility.** In the previous subsection, we sketched a basic rationale for the existence of utility functions, as part of an evolutionary mechanism to shift partial control from evolution to the individual in order to exploit the individual's familiarity with idiosyncratic current circumstances. We turn now to a more detailed account of the implementation of utility in the brain. We

---

[9] These offspring numbers are only meant capture the issue in a simple way.

argue that neural processes represent utility in a concrete and quantifiable fashion and that this drives choice rather than being merely a ghostly reflection of it.

Robson and Samuelson (2024) propose a three way nested classification of possible notions of utility—

1) Representational utility. This is the standard view that utility is merely an analytical device used to describe behavior, but without any independent existence. Utility is then described as "preferences" to emphasize the weakness of the notion. The culmination of this notion is the theory of revealed preference, which involves a set of axioms that constitute the irreducible nub of demand theory, so that further simplification must have observable consequences for demand.[10]

2) Substantive utility. Robson and Samuelson (2024) use this term to denote utility that has a concrete existence, as neural activity in the brain, and is the form of utility advocated here. Further, utility does not merely reflect choice but drives it and so it is representational and more. The payback from adopting this perspective is that it suggests properties of utility that would not be as compelling under revealed preference. Once utility is seen as neural activity, in particular, it is natural to consider how the prevalence of neural noise leads to possibly erroneous choice, making adaptation a central feature.[11] We sketch these implications here; they are developed further by Robson and Samuelson (2024).

3) Hedonic utility. In a more tentative spirit, Robson and Samuelson suggest that utility is not merely substantive but can be interpreted in terms of affect or pleasure. This notion of utility has a lineage stretching back to Bentham (1798) who advocated a notion of cardinal and hedonic utility. Bentham was mainly preoccupied with social welfare. In order to make social trade-offs, utility must be cardinally comparable in some fashion across individuals; Bentham thought the hedonic element would provide a basis for this. The modern revealed preference approach was the culmination of the reaction against such earlier views, a reaction

---

[10] Revealed preference theory is outlined in Mas-Colell, Winston and Green (1995, Chapter 1).

[11] Glimcher (2011) advocates a related concept of utility as generated by neural processes and subject to inescapable error.

that came to be organized around “Occam’s Razor”’—the principle that redundant hypotheses should be dropped.[12] An hedonic view of utility is involved in modern discussions of “life satisfaction”—see Clark, Flèche, Layard, Powdthavee and Ward (2018) , for example.

We turn now to evidence that buttresses the notion of utility as substantive at least and derive the implications. We briefly mention the more incomplete evidence bearing on hedonic utility.

Stauffer, Lak and Schultz (2014) (SLS) provide evidence that utility is directly represented in the brain and apparently drives choice. The evidence concerns primates and some of it involves direct observation of neural activity.

Two monkeys, A and B, are offered juice rewards. By finding the certainty equivalents to various binary gambles, SLS first derive von Neumann-Morgenstern (vNM) utility for each monkey, using a familiar procedure that is purely behavioral. As might be found for human subjects using a similar procedure, the resulting vNM utility function is S-shaped. SLS check that this vNM utility function predicts choices between the certainty equivalent and various out-of-sample gambles. This construction is represented in Figure 1.

They now show that vNM utility is realized in the brain by dopamine-producing neurons. Most basically, perhaps, a burst of activity of the dopamine neurons is associated with an unanticipated physical reward. Indeed, these unanticipated neural firing rates are quantitatively similar to behavioral vNM utility. See Figure 2.

The precise role of dopamine in shaping behavior remains controversial, despite a backdrop consensus that dopamine has some connection to both behavior and emotion. On the one hand, Schultz (2015), for example, stresses that dopamine plays a central role in inducing learning. On the other hand, Berridge (2007, 2009) is a leading proponent of the view that dopamine reflects the desire for something, and is not tied rigidly to learning. He buttresses his argument by pointing to instances where learning occurs without the aid of dopamine.

---

[12] Kahneman, Wakker, and Sarin (1997) is a modern revival of some of Bentham’s views.

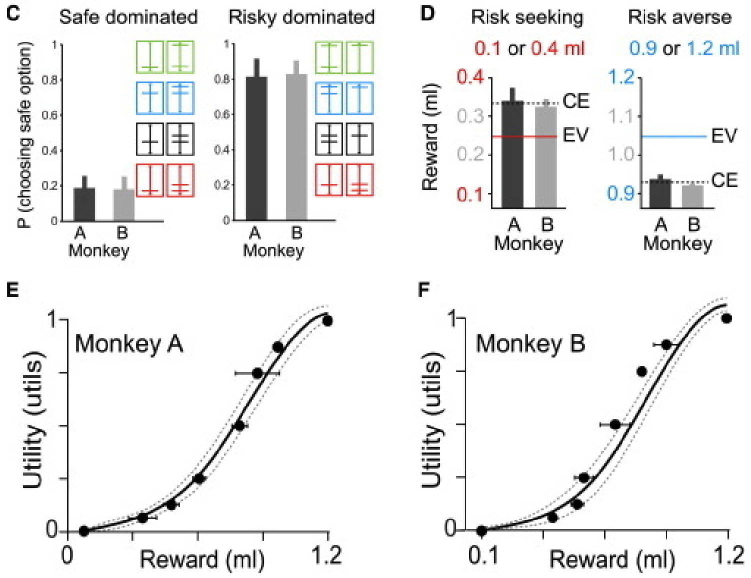

FIGURE 1. This describes the purely behavioral construction of von Neumann Morgenstern (vNM) utility for each of two monkeys presented with a choice between a binary 50-50 gamble over juice rewards and a certain outcome. This experiment requires a huge number of repetitions in order that that the monkeys appreciate what the gambles involve. This figure excerpts Stauffer, Lak, and Schultz's (2014) Figure 1. Used under Creative Commons.

A typical current view is that pleasure often arises from a few small neural sites that are spatially distributed in the brain. These sites involve so-called opioid receptors. Recent progress (e.g., Kragel *et al.* (2023)) has been made in recognizing a characteristic pattern of activation that can be associated with pleasure, where such pleasure summarizes multiattribute rewards. Berridge and Kringelbach (2015) provide a recent survey of research on the neural foundation and function of pleasure. They describe the similarity of the circuitry switched on in diverse pleasures ranging widely from food to music. Such a system apparently uses a common currency of pleasure, permitting a flexible response to novel stimuli, so a wide variety of options can be evaluated, compared and chosen among.

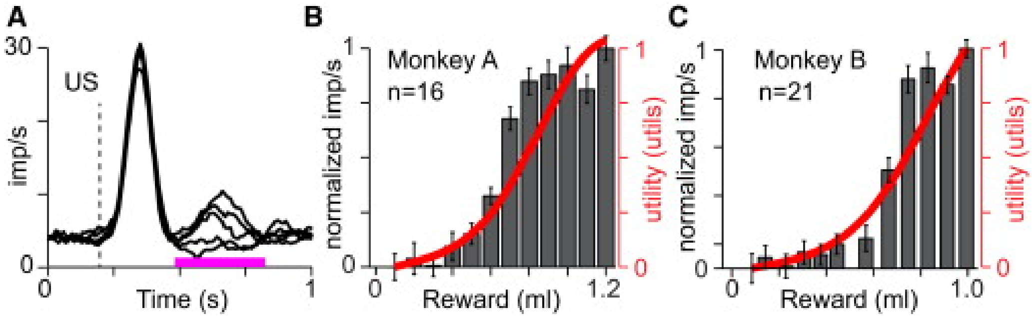

FIGURE 2. This reports the results of presenting each subject monkey with various, unanticipated quantities of the reward. The bar graph depicts the firing rates of dopamine neurons in response to the reward; the red line is the vNM utility obtained behaviorally as in Figure 1. The firing rates used derive from the region in the first panel that is underlined, which is where the effect of the differing amounts of juice is seen. This is Stauffer, Lak, and Schultz's (2014) Figure 4. Used under Creative Commons.

The consumption of food is an illuminating example of how behavior may be driven by hedonic utility, as represented by opioid receptor activity. Tuulari *et al.* (2017) find experimentally that food causes a significant surge in endogenous opioids in humans, after consumption of either a palatable or a non-palatable meal.

Altogether, a tentative picture that emerges for the mechanism behind feeding behavior is that (i) dopamine either helps control learning about cues leading to food acquisition (as described by Stauffer, Lak and Schultz (2014)) or orchestrates "wanting" food (as hypothesized by Berridge (2009), for example), or both, while (ii) opioid receptors, for example, generate a pleasurable response to the actual consumption of food.

With respect to either dopamine or opioid receptors the evolutionary history is deep, going back well beyond hunting and gathering, to ancestral non-human species. Dopamine has a huge variety of functions, many of which appear unrelated to the present focus. Costa and Schoenbaum (2022) provide a summary account of both the functions of dopamine and its evolutionary antiquity, as evidenced by its being currently found in fruit flies and sea snails. Dreborg *et al.* (2008), trace opioid receptors back at least 450 million years to the first vertebrates with jaws. However, it is plausible that hunting and gathering refined the mechanisms implementing utility.

Freed [36, Chapter 11] is a compelling and clear textbook treatment of current knowledge of the neuroscientific basis of motivation and choice. He provides an engaging account of the roles of dopamine and opioid receptors, in particular.

2.2. **A Model of Adaptive Substantive Utility.**

The previous subsection outlined evidence for noisy and adaptive substantive utility. We now consider two models of how substantive utility is realized in the brain, the first of these is in the limelight in Sections 2.2.1-2.2.6 and the second is sketched in Section 2.3. These models share a common motivation and differ only in the details as to how they capture this.

As a simile, consider the issue of measuring a voltage in physics. With many voltmeters, it is necessary to first guess the right range in which the unknown voltage lies. If you guess a range that is too high, the needle might not budge, and so no useful information is obtained; if you guess a range that is too low, the needle will peg against the stop at the top, or the meter will blow out, and again no information can be obtained.

Analogously, we propose that bundles of goods are only perceived with imprecision, and that this imprecision leads to errors in choice that can be offset by adaptation of perception. In Section 2.1 we summarized evidence that von Neumann Morgenstern utilities correspond to measurable events in the brain, firing rates of dopamine neurons, in particular. Given that utility is concrete and noisy like this, there will be an evolutionary advantage in designing utility functions that shift with the circumstances. That is, it will pay to have more precise perception where the traffic is heavy and less precise perception where the traffic is light. If traffic patterns shift, then it will pay for perception to shift in response. There is an advantage in putting the "steep part" of the utility function where it counts.

FIGURE 3. Voltmeter with various ranges for voltages. Creative Commons. A M Felicisimo.

2.2.1. *A Threshold Model.* This basic threshold model here derives from Robson (2001) which is much elaborated in Robson, Whitehead, and Robalino (2023).

Suppose each of two alternatives, denoted by $i \in \{1,2\}$, generates an outcome $y_i \in [0,1]$ that is independently drawn from the continuous distribution function

F with density f. We assume that F has values from $[0,1]$. Evolution would prefer the agent to secure the larger realization y_i.

The set of outcomes $[0,1]$ is divided into $N+1$ intervals by N thresholds $0 < x_1 < .... < x_N < 1$, where we take $x_0 = 0$ and $x_{N+1} = 1$. Evolution's design problem is then to choose the threshold values of $\{x_1, \ldots x_N\}$.[13]

The individual characterized by the thresholds $0 < x_1 < .... < x_N < 1$ observes only the interval to which the realization y_i belongs. Faced with two realizations in different intervals, the individual can then be sure the higher realization is that falling into the higher interval. If the two realizations lie in the same interval, then the alternatives are indistinguishable and the individual makes a mistake with probability $1/2$.

2.2.2. *Minimizing Error.* Suppose that evolution's goal is to induce behavior from the agent that minimizes the probability that the agent chooses the alternative with the lower realized value y_i. First, suppose evolution can choose only a single threshold, denoted by x. The probability of making an error is

$$\begin{aligned} PE(1) &= (1/2)\Pr\{y_1, y_2 < x\} + (1/2)\Pr\{y_1, y_2 > x\} \\ &= (1/2)(F(x))^2 + (1/2)(1-F(x))^2 \\ &= (1/2)\xi^2 + (1/2)(1-\xi)^2, \end{aligned}$$

where y_1 and y_2 are the outcomes of the two arms and $\xi = F(x)$. That is, an error occurs if and only if y_1 and y_2 are either both below or both above x, when the probability of an erroneous decision is $1/2$.

This objective function is a convex function of ξ, so the first-order condition is sufficient—

$$\frac{dPE(1)}{d\xi} = \xi - (1-\xi) = 0,$$

and hence $\xi = F(x) = \frac{1}{2}$. Thus x should be the median of the distribution described by F.

---

[13] Evolution must also choose the number of thresholds, N, given the costs and benefits of N. Whatever N is chosen, however, the thresholds should be optimally allocated.

Evolution's preferred utility function then depends on the distribution F generating realizations.

What if the agent is able to identify multiple thresholds? Suppose there are N thresholds

$$0 = x_0 < x_1 < x_2 < ... < x_N < x_{N+1} = 1.$$

The probability of error is now

$$\begin{aligned} PE(N) &= (1/2)(F(x_1))^2 + ... + (1/2)(F(x_{n+1}) - F(x_n))^2 + ... + (1/2)(1 - F(x_N))^2 \\ &= (1/2)(\xi_1)^2 + ... + (1/2)(\xi_{n+1} - \xi_n)^2 + ... + (1/2)(1 - \xi_N)^2, \end{aligned}$$

where $\xi_n = F(x_n)$ for $n = 1, ..., N$. Each term is the probability that y_1 and y_2 both fall between the same pair of threshold values, times the probability $1/2$ that an erroneous decision then occurs.

It follows readily that the optimal allocation of thresholds, the one that minimizes $PE(N)$, is given by

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

The thresholds that minimize the probability of error are thus equally spaced in terms of probability. If $N = 9$, the thresholds should be at the deciles of the distribution.

The probability of error $PE(N)$, when the thresholds are (optimally) evenly spaced in terms of probability, is

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

Each term is the probability that y_1 and y_2 both fall between a common pair of thresholds, multiplied by the probability of one half of then making an error. Additional thresholds are advantageous—evolution prefers larger values of N. Increasing N sharpens the agent's perceptual precision. The optimal number of thresholds balances the benefits of better decision making against the cost of more precise perception.

In the limit as N gets arbitrarily large, the limiting density of thresholds is the density function f. The simplest utility function, which corresponds to counting thresholds in the finite N case, would assign the realization y to y's percentile rank in the distribution of thresholds. For a given distribution f of thresholds, the utility of realization y is then

$$U(y) = \int_0^y f(z)dz.$$

Having an infinite number of thresholds ensures that the agent never makes a mistake. We appeal to this limiting case mainly to shed light on the shape of the utility function. This shape remains, even in the limit when there are no mistakes, serving to illustrate what will happen for large finite N.

Robson, Whitehead and Robalino (2023) discuss the important alternative that the thresholds are chosen to maximize the fitness of the alternative chosen. Although maximizing fitness is an appealing criterion, we confine ourselves here to the more tractable case of minimizing the probability of error.

2.2.3. *Dynamic Adaptation.* The implication of the previous discussion is that evolution would like to tailor the utility thresholds to the distribution F from which outcomes are drawn. Intuitively, the thresholds are optimally concentrated "where the heavy traffic is." In regions where the density f is large and hence realizations are likely, thresholds should be relatively prevalent, giving a utility function that is relatively steep and hence capable of finely distinguishing alternatives. In regions where the density f is small and hence realizations rare, thresholds are sparse, giving a relatively flat utility function that permits more mistakes, conditional on being in these regions in the first place.

This generates a need for utility to adapt rapidly. Consider then the following simple adaptive learning process. We focus here on the case of the probability of error criterion.[14] For a fixed distribution F, this learning process converges to an approximation of the optimal configuration of thresholds. The process is simple

---

[14] See Robson, Whitehead and Robalino (2023) for a detailed treatment of a generalized model that treats the maximizing fitness case in a limiting sense as $N \to \infty$.

and requires no parametric information about F. As long as this adjustment process can proceed relatively rapidly, compared to the rate at which the environment fluctuates, the agent will most of the time employ a utility function that is nearly optimal.

Suppose that the thresholds must be chosen on a finite grid $\mathcal{G} = \{0, \varepsilon, 2\varepsilon, ..., G\varepsilon, 1\}$, for an integer G such that $(G+1)\varepsilon = 1$. Evolution must choose N elements from this set to comprise the thresholds in the individual's choice rule. It is technically convenient to have a finite number of possible thresholds. We will consider the limiting case as ε approaches zero.

We let $x_n^t \in \mathcal{G}$, where $0 < x_1^t < ....x_N^t < 1$, at time $t = 1, 2, ...$, denote the various thresholds at time t. Consider the following rule for adjusting the thresholds:

Suppose the period t outcome y_t is in $[x_{n-1}^t, x_n^t)$. Then—

If x_{n-1}^t is not $x_0 = 0$, then $x_{n-1}^{t+1} = x_{n-1}^t + \varepsilon$, and if x_n^t is not $x_{N+1} = 1$, then $x_n^{t+1} = x_n^t - \varepsilon$.

The basic motivation for the rule is that intervals shrink when realizations land in them. This ensures that, in regions where outcomes are relatively likely, the intervals are small, allowing the agent to discriminate finely between outcomes that occur there.

In response to this learning rule, the values of the thresholds will constantly fluctuate. This fluctuation will continue even if the distribution F from which realizations are drawn is permanently fixed. These fluctuations have the upside of creating the ability to adjust to changes in the distribution F.

More formally, this adjustment process induces a Markov chain whose state space consists of vectors of thresholds. The Markov chain is ″irreducible″—so that there is a time T such that, given any initial vector of thresholds, there is a positive probability of reaching any other vector at T. It follows that there is a unique ″invariant distribution″ and that the distribution of state vectors converges to this from any initial state.

In the limit with small ε, the thresholds adjust so as to minimize the probability of error. That is, Robson, Whitehead and Robalino [2023, Theorem 3.1] prove:

Proposition Fix N. In the limit as $\varepsilon \to 0$, the invariant distribution of the nth threshold x_n^t converges to a point mass at x_n^*, where $F(x_n^*) = n/(N+1)$, for $n = 1, ..., N$.

Hence, as the grid of possible values becomes finer, the adjustment rule generates thresholds arbitrarily close to the equally-spaced-in-terms-of-probability rule that is optimal for minimizing the probability of error.

We do not claim that this learning *process* is optimal. If the distributions F the agent encounters are drawn from some parametric class, then the efficient mechanisms should exploit this restriction. To talk about optimal learning, in general, one would specify the stochastic process governing the evolution of the distribution F, equip the agent with Bayesian beliefs over the identity of the current distribution, and solve a discounted expected utility maximization problem. This might be beyond evolution's power to implement.

The current exercise confirms just that there exists a simple process leading the agents to an ultimately approximately optimal response to any fixed distribution F. There are undoubtedly many other ultimately approximately optimal behavioral rules.

2.2.4. *Preference for Growth.* Consider now a dynamic application of adaptive utility, involving preferences over income profiles. The motivating empirical background here is a clear tendency for wages to increase with a worker's tenure at a firm. For example, Dustmann and Meghir (2005) provide evidence of the returns to tenure, finding that skilled workers enjoy substantial wage gains with experience, with strong gains from staying with the same firm. Unskilled workers' wages grow even more strongly early in their tenure.[15]

A well-known explanation (for example, Becker (1964)) for rising wage profiles is that firms invest in workers early in their careers, and then pay wages that increase with tenure as an inducement for the workers to remain with the firm, allowing the firm to reap the return on its investment. This investment may take the form

---

[15] This material on a preference for growth is drawn from Robson, Whitehead and Robalino (2023), which also presents simulations. See also Robson and Samuelson (2025).

of explicit training or education, but may also be implicit—if worker productivity increases with tenure, no matter what the reason, then the early years of employment are effectively an investment, with the subsequent higher wages reflecting the importance of retaining workers who have longer tenures and are hence more productive.

Frank and Hutchens (1993), on the other hand, argue persuasively that conventional preferences are unable to fully explain the ubiquity of increasing remuneration streams. They show that intercity bus drivers and commercial airline pilots both earn wages that increase with tenure. They argue that in the case of airline pilots, most of their training consists of prior experience in the military, so that there is little explicit investment to be recouped by an airline. In the case of intercity bus drivers, there seems to be little room for improved performance with a long tenure. Frank and Hutchens then argue that increased wage profiles should be partly explained by including a primitive preference for a growing rate of consumption in the utility function. We propose an evolutionary derivation of such preferences.

Let us consider a scenario in which an agent faces a steadily increasing income stream. It can be shown that the agent's utility evaluations will adapt, converging to a constant level arising from the sequence of increasing consumption. Furthermore, this constant level is raised by the increasing nature of the sequence.

If the agent makes choices among such increasing (or decreasing) income streams, we assume that she correctly anticipates how her utility function will evolve with the consumption stream, and evaluates a consumption stream according to its induced average utility.

Take the model of Section 2.2.1, with N thresholds, whose values are adjusted according to the process introduced in Section 2.2.3. To capture a trend in incomes, we fix a distribution F^0, and suppose that rewards at time $t \geq 0$ are drawn from the distribution F^t, satisfying

$$F^t(x) = F^0(x - \alpha t).$$

Hence, the distribution from which rewards are drawn shifts at a constant rate α. The distribution F^0 has support on a unit interval $[x_0, x_0 + 1]$, while F^t has support on $[x_0 + \alpha t, x_0 + \alpha t + 1]$. We allow the case of $\alpha > 0$, so that the distribution is shifting upwards, as well as $\alpha < 0$ and a downward shifting distribution.

Let y_n^t denote the value of the nth threshold at time t. Define

$$\hat{y}_n^t = y_n^t - \alpha t.$$

We refer to $\hat{y}_n^t$ as the detrended value of the nth threshold at time t, and refer to y_n^t as a raw threshold.

The distribution of the raw thresholds y^t will steadily increase or decrease in response to the shift in F^t. In contrast, the detrended thresholds converge within the unit interval. That is, the argument of Robson, Whitehead and Robalino [2023, Lemma 7.1] shows:

Proposition Fix the number of thresholds N. For sufficiently small α, there exists an interior invariant distribution of the detrended thresholds $\hat{y}_n^t$. In the limit as $\varepsilon \to 0$, the invariant distribution of detrended thresholds converges to the distribution that assigns a point mass to the vector with components $\hat{y}_n^*$, $n = 1, ..., N$ that are the unique solutions to

$$(F^0(\hat{y}_{n+1}^*) - F^0(\hat{y}_n^*)) = (F^0(\hat{y}_n^*) - F^0(\hat{y}_{n-1}^*)) + \alpha \tag{2.1}$$

for $n = 1, ..., N$.

This implies that average utility is raised by increasing realizations and lowered by falling realizations. The intuition is that with increasing realizations, the thresholds are constantly lagging behind the distribution of realizations. This lagging behind means that the raw thresholds y_n^* are being constantly tugged upwards, as they need to be. On average, the realizations are then good news. With decreasing realizations, the thresholds are again constantly lagging behind the realizations, but this time the realizations are, on average, bad news.

Robson and Samuelson (2025) examine a more dramatic version of this phenomenon, for an adjustment rule that does well, in a limiting sense, in terms of expected fitness rather than the probability of error, which we focussed on in this chapter. They show that people might prefer *time-varying* income streams over comparable constant streams. The setting itself is quite special, and the result holds for a specific adjustment rule. The example nevertheless illustrates how a mechanism that is only approximately optimal in typical situations might generate startling choices in rather different circumstances.

2.2.5. *Hedonic Treadmill.* We discussed how a preference for rising income might arise from adaptive utility. The evidence in favor of adaptive utility is compelling—both evidence concerning its neuroscientific basis and evidence of its predicted effects. A less clear picture emerges for hedonic utility, which is not merely substantive but associated with reported pleasure or happiness. In a more tentative spirit, then, we turn to the most well-known of phenomena that concern adaptive pleasure—the so-called "hedonic treadmill".

David Schkade and Daniel Kahneman (1998) found that students in Ann Arbor and Los Angeles reported similar levels of life satisfaction. However, students in Ann Arbor predicted significantly higher levels of life satisfaction for those living in LA. Schkade and Kahneman describe this as a "focusing illusion". For them, there is a distinction between "decision utility"—which would be the basis of a decision to move to LA—and "experienced utility"—what you actually experience once there, and the "hedonic honeymoon" is over.

It is not surprising somehow that students in Michigan would predict greater life satisfaction for those living in LA or for themselves if they would move there. Suppose, for example, you compare surfing in the winter in Michigan...

FIGURE 4. Surfing in the winter on Lake Michigan. Image generated by Google Gemini.

...with surfing in the winter in California...

FIGURE 5. Surfing in the winter in Encinitas, California. Image generated by Google Gemini.

Since realized long term life satisfaction in LA does not match this rosy prediction, Schkade and Kahneman assert that decision utility is then *wrong*—that you might well make a mistake in deciding to move to LA.

Robson and Samuelson (2010) use the Rayo and Becker (2005) model to generate distinct decision and experienced utilities, so agreeing with Schkade and Kahneman up to a point, but crucially argue that there is no suboptimality here. That is, decision utility is the basis of the best possible choice about moving, and experienced utility the basis of the best possible decisions once in LA. That is, far from being a problem, adaptation represents the biologically optimal response to limited perceptual accuracy.

2.2.6. *Treadmill—Simulations.* Reconsider the preference for growth model of Section 2.2.4 with the parametric class of cumulative distribution functions $F(x) = x^{\gamma}$ and probability density functions $f(x) = \gamma x^{\gamma-1}$, for all $x \in [0, 1]$, with parameter $\gamma > 0$. Suppose the grid size is $\varepsilon = 0.0005$. Set $N = 9$ and take $100,000$ periods, where $\gamma = 1$ for the first $20,000$ periods and $\gamma = 5$ thereafter up to $100,000$.

Figure 6 demonstrates how, assuming that substantive utility induces experienced and reported pleasure, the adaptive utility model illuminates the phenomenon of the hedonic treadmill.

2.2.7. *Variants.* Variants of the model can also shed light on phenomena that directly involve variation across individuals rather than across the outcomes of a gamble. This could involve hypothesizing that the choices that an individual needs to make scale with the distribution of wealth or income.

Consider, for example, the Easterlin paradox, as described originally in Easterlin (1974) , which is the observation that, although reported life satisfaction typically increases within a nation, or across nations, with greater income, there is less evidence of a secular rise in life satisfaction with generally rising incomes.[16]

Consider the hedonic treadmill in this light. There is steady progress—a positive secular trend—in the distribution of income and hence steady progress in the

---

[16] The Easterlin paradox generated a vast literature. Easterlin and O'Connor (2022) is a recent paper that addresses some of the criticisms it has attracted.

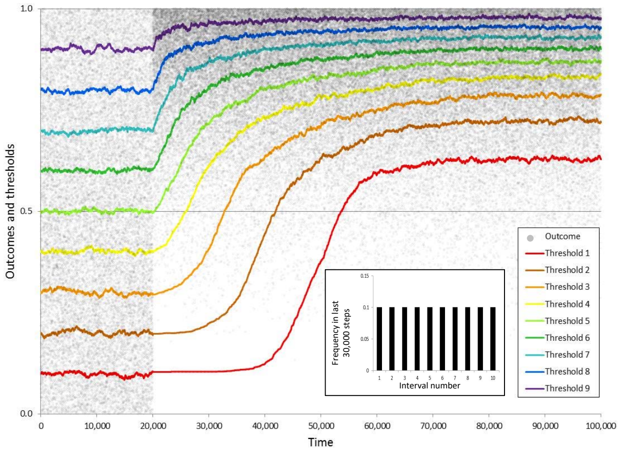

FIGURE 6. Simulations of the hedonic treadmill. The thresholds move to the new deciles, as shown by the ultimately equal probabilities of each interval for the new distribution. Average utility rises sharply at first, with the shift to a first-order stochastically dominant distribution, but then subsides to its original level, despite the permanence of the shift.

outcomes of the relevant decisions faced by an individual. The growth process induces a higher level of happiness but not happiness that grows over time. Contemporaneously, happiness is higher for higher outcomes, as is all consistent with Easterlin. The hedonic treadmill model presumes that the decisions faced involve purely contemporaneous tradeoffs, which might be a plausible approximation.

A further related application would be to the concern with status, especially the life satisfaction deriving from that. Such a concern is often associated with Veblen (2017). Duesenberry (1949), for example, showed that the fraction of income that is consumed depends largely on one's relative position in the wealth distribution. That is, for example, someone at the lower quintile, say, of the wealth distribution currently consumes a similar fraction of income to someone at the same relative position 20 years ago and not someone with the same real income, who would then have been relatively richer.

Robson (1992) argues that hypothesizing a concern with status in the sense of the position in the wealth distribution accounts for the observations made by Friedman and Savage (1948) concerning the taking and avoidance of risk. That is, the avoidance of gambles involving moderate gains and losses, but the taking of lotteries involving small losses with high probabilities and large gains with small probabilities suggests that von Neumann-Morgenstern utility might be concave at first but then become convex. If this utility were a function of wealth alone, the large gains with small probability should decrease with growth, contrary to observation. Hypothesizing an additional concern with relative position can account for large gains that increase with growth, as is actually observed.[17]

2.3. **Encoding and Decoding Noisy Signals.**

The threshold model that is the main focus here arranges outcomes into bins, with utility constant within a bin and taking jumps from one bin to another. We could view this as a simple model of a just-noticeable difference rather than a literal description of utility functions. Netzer, Robson, Steiner and Kocourek (2024) (NRSK) present a model of noisy perception that smooths some of the lumpiness of the threshold model.

The model of NRSK endows the individual with a continuous set of perceived outcomes. The individual perceives the realized value of each outcome, transformed by a nonlinear encoding function and perturbed by an added noise term. For tractability, NRSK focus on the limiting case where there are many observations on each outcome. If the individual has a correct view of the statistical relationship among the various outcomes, the individual will draw the appropriate inferences about each observation and will then make appropriate risk-neutral choices.

Matters become more interesting when the individual misperceives this statistical relationship, perhaps mistakenly coalescing outcomes. Such misperceptions might arise because it is beyond the power of evolution to endow the agent with a perfect understanding of the environment. In this case, even with a large number

---

[17] See Ray and Robson (2010) for a dynamic version of this model.

of observations, individuals will choose in a way that reflects the nonlinear encoding function. If there are just two outcomes, for example, which the individual believes must have the same realization, the individual evaluates the binary gamble in terms of the expectation of the encoding function, where the expectation reflects the sampling rates on the two outcomes.

This model avoids awkward conceptual issues concerning the sharp edges of the threshold model. That is, at each threshold there is a discontinuous jump in the signal the individual obtains. On the other hand, the continuous model is less tractable. It is more difficult to analyze even the static case, but the difficulties are exacerbated in the dynamic case. The threshold model allows one to describe the process of dynamic adjustment in terms of how a finite number of thresholds shift with experience. In a continuous model, on the other hand, adjustment concerns the much less tractable problem of adaptation of an encoding *function*.

In the current section, we considered utility as instantiated by neural activity in the brain. We circled the wagons most tightly about this notion of utility as substantive, where such neural activity drives choice. The inevitable neural noise then makes adaptation desirable and we derived the observable consequences. We also considered more tentatively the implications for hedonic utility—so reported pleasure or life satisfaction would exhibit adaptation. We now switch gears, adopting an approach founded on a literally evolutionary perspective. We take a theoretical view of how different time preferences and different attitudes to risk would have different long run demographic success, thus selecting favored time preferences and risk attitudes. We investigate the empirical implications.

3. Evolution of Attitudes to Idiosyncratic Risk and Time Preference

Two basic attributes of preferences are attitudes to risk and time preference. It is also quite plausible that these two attributes were shaped by natural selection in hunter-gatherer societies. A von Neumann-Morgenstern utility function with constant relative risk aversion has properties that simply scale risk attitudes from an impoverished evolutionary context to a better-off modern context. Furthermore, time preference has a dimensionless quality that could have retained its precise meaning in widely different circumstances.

We first consider the straightforward fashion in which evolution shapes attitudes to idiosyncratic risk. This straightforward attitude derives from the biological criterion for success in this setting—expected offspring. This is a building block in all that follows—including, for example, results on time discounting, as well as the results concerning aggregate risk presented in Section 4. We show that attitudes to idiosyncratic risk might have evolved to reflect the environment, to be "plastic". That is, an evolutionary approach does not force everyone into a straitjacket of identical preferences.

We then consider how the pure rate of time preference might have evolved in a hunter-gatherer society. We consider how the receipt of resources at different ages would have a differential effect on biological success. This demographic approach derives time preference from age rather than from time into the future, which distinguishes this approach from the conventional one. Even when this demographic approach generates rates of time preference that vary with age, this does not imply preference reversals. We show how time preference may also exhibit plasticity with a time-varying environment and how it might reflect geographic variation in evolutionary history.

3.1. **Attitudes to Idiosyncratic Risk.** Consider a population that reproduces parthenogenetically, that is, without sex. Each of individuals $n = 1, ..., N$ has a distribution of offspring given by an independent draw $\tilde{x}_n$ of the random variable $\tilde{x}$ with mean $\bar{x}$.[18] It follows that the growth rate of the population is

$$\frac{\sum_n \tilde{x}_n}{N} \to \bar{x},$$

as $N \to \infty$, with probability 1, by the strong law of large numbers.

That is, individuals evaluate options in terms of the expected offspring they generate. For example, suppose that resources $r \geq 0$ translate into expected offspring according to the function $x = f(r) \geq 0$, where the random variable representing offspring conditional on r is independent across individuals. Then, gambles over

---

[18] That is, the probability of $i \geq 0$ offspring under $\tilde{x}$ is $p_i \geq 0$ for some p_i where $\sum_i p_i = 1$.

r, given as the random variable $\tilde{r}$, that are also independent across individuals are evolutionarily evaluated as $E(f(\tilde{r}))$. That is, they are evaluated in terms of expected utility where the role of von Neumann-Morgenstern utility is played by the function f that connects resources to expected offspring.

3.1.1. *Plasticity of Attitudes to Idiosyncratic Risk.* We now illustrate the general point that it is generally advantageous for behavior to be be "plastic"—to respond to the environment—showing that evolution is readily capable of generating a spectrum of behavior, adapted to time varying circumstances.

We first consider attitudes to risk. Suppose then that the expected offspring generated by a random allocation of resources $\tilde{r}$ depends on the environment ξ as $Ef(\tilde{r}, \xi)$, say. Suppose ξ has two possible realizations ξ_1 and ξ_2.[19] Conditional on ξ_i expected offspring for a particular gamble $\tilde{r}_1$ exceeds that for another gamble $\tilde{r}_2$, but conditional on ξ_2, the reverse is true. That is,

$$Ef(\tilde{r}_1, \xi_1) > Ef(\tilde{r}_2, \xi_1)$$

but

$$Ef(\tilde{r}_2, \xi_2) > Ef(\tilde{r}_1, \xi_2).$$

Suppose that the environment is independently drawn each period , with the probability of state ξ_i being p_i for $i = 1, 2$, where $p_1 + p_2 = 1$.[20]

Suppose the large total population at date T is now given by $P(T)$, where $P(0) = 1$. If there is no information about the state, the gamble chosen cannot depend on the state. If gamble $\tilde{r}_i$ is taken throughout, then the population at date T is

$$P(T) = (Ef(\tilde{r}_i, \xi_1))^{T_1} \, (Ef(\tilde{r}_i, \xi_2))^{T-T_1}$$

where T_1 is the random variable giving the number times that state ξ_1 is in effect, so that $T - T_1$ is the number of times ξ_2 is in effect. The limiting growth rate of

---

[19] A particular case might be that the individual becomes more or less risk averse with the change in the environment.

[20] The iid nature of the environment can be readily significantly generalized. For example, it could be a Markov chain with limiting probabilities p_1 and p_2.

population this choice of gamble induces is then

$$\frac{\ln P(T)}{T} = \frac{T_1}{T} \ln \left(Ef(\tilde{r}_i, \xi_1) \right) + \frac{T - T_1}{T} \left(Ef(\tilde{r}_i, \xi_2) \right) \to p \ln \left(Ef(\tilde{r}_i, \xi_1) \right) + (1-p) \left(Ef(\tilde{r}_i, \xi_2) \right) = R_i,$$

say, with probability one, as $T \to \infty$, by the strong law of large numbers. This limiting growth rate characterizes the evolutionary success of a type that takes gamble r_i always.

These calculations consider aggregate uncertainty concerning the environment $\xi_i; i = 1, 2$, and presage the treatment of attitudes to aggregate risk in Section 4. The choices available here are different idiosyncratic gambles *conditional* on the aggregate state $\xi_i, i = 1, 2$. The focus in Section 4 is to compare gambles that involve aggregate risk with those that involve idiosyncratic risk. The most general gamble there involves a random aggregate state with i.i.d gambles conditional on each state. Choice among these gambles expresses a differential concern with aggregate risk.

Suppose now the shift in the environment is signaled so that an individual can condition on the signal. The type that will do best will then maximize expected offspring conditional on the environment. That is, the total population at T is now

$$P(T) = \left(Ef(\tilde{r}_1, \xi_1) \right)^{T_1} \left(Ef(\tilde{r}_2, \xi_2) \right)^{T - T_1}$$

with a limiting growth rate

$$p \ln \left(Ef(\tilde{r}_i, \xi_1) \right) + (1-p) \left(Ef(\tilde{r}_i, \xi_2) \right) = R^*,$$

say, where it follows immediately that $R^* > R_i$, $i = 1, 2$. Hence the type that adapts to the environment outperforms either type that does not, eventually dominating the population.

This argument demonstrates the advantage of behavioral plasticity, where the optimal choice to be made is conditional on the environment, and where such choice can be made rapidly. There might be an underlying genetically controlled predisposition to detect and apply information that induces such rapid adaptation, but the predisposition itself might have arisen by the much slower process of natural selection, given repeated exposure to the various environments.

Such behavioral plasticity is a familiar phenomenon in biology. For example, female North American red squirrels responded to the abundance of white spruce cones created by global warming by advancing their average age at the first birth of offspring by 18 days (see Réale, McAdam, Boutin, and Berteaux (2003)). Such plasticity is not confined to behavior. For example, many species of turtle change sex depending on the temperature of the water (see Bull (1980), for example).

3.2. **Time Preference.** [21]

We sketch now an account of the evolution of the rate of time preference, based on the demographic variables of survival to each age and fertility at each age.[22]

For simplicity, although the results here hold more generally, agents live three periods, producing x_i expected offspring at ages $i = 1, 2, 3$. Three periods is the minimum life span needed to consider any variation in the rate of time preference with age. However, most of the results found here extend straightforwardly to more periods.

The probability of survival from age i to age $i+1$ is S_i for $i = 0, 1, 2$. Suppose there is a large population described by the vector $N_a(t)$ for ages $a = 1, 2, 3$ and dates $t = 0, 1, ...$ Given that survival and fertility are independent risks, this population evolves deterministically as

$$N_1(t+1) = S_0\left(N_1(t)x_1 + N_2(t)x_2 + N_3(t)x_3\right).$$

That is, there are $N_1(t)x_1 + N_2(t)x_2 + N_3(t)x_3$ newborns at t given the numbers and fertility of the three adult ages and these newborns survive with probability S_0 to $t+1$.[23] The evolution of the older ages is very simple. Each cohort comes from the cohort one period ago that was then one period younger, allowing for interim mortality—

$$N_2(t+1) = S_1 N_1(t), \tag{3.1}$$

$$N_3(t+1) = S_2 N_2(t), \tag{3.2}$$

---

[21] A fuller treatment of the material in Sections 3.2-3.2.2 is in Robson and Samuelson (2024).

[22] This material is drawn from Robson and Samuelson (2007).

[23] The notation suppresses, explicitly, the number of newborns.

for $t = 0, 1, ...$ Using (3.1) and (3.2) for $N_2(t)$ and $N_3(t)$,

$$N_1(t+1) = S_0 N_1(t) x_1 + S_0 S_1 N_1(t-1) x_2 + S_0 S_1 S_2 N_1(t-2) x_3, \text{ for } t = 2, ... \tag{3.3}$$

In steady state growth, the proportions of each age class remain constant over time, so that each age class grows with the same factor. If there is steady state growth with growth factor λ, then, in particular, $N_1(t+1) = \lambda N_1(t)$, for all t. It follows from (3.3) for $N_1(t+1)$ that

$$1 = \frac{S_0 x_1}{\lambda} + \frac{S_0 S_1 x_2}{\lambda^2} + \frac{S_0 S_1 S_2 x_3}{\lambda^3}, \tag{3.4}$$

which is the Euler-Lotka equation. This equation determines λ from the discounted value of expected offspring at each adult age, $S_0 S_1 x_2$ at age 2, for example. The discount factor λ is raised to a power corresponding to age.

Although this equation cannot be solved in closed form, whenever the maximum age is more than a few periods, it is readily seen to always have a unique real positive solution for the growth factor λ.

Perron-Frobenius Theorem. Given any initial population vector, the population settles down into steady state growth with growth factor given by the unique positive solution of the Euler-Lotka equation.[24]

The main content of this theorem is convergence to steady-state growth from arbitrary initial conditions. That is, once that is given, it is generally straightforward to derive the Euler-Lotka equation, basically as we did in this example.

Consider now the set of the x_i that yield the same value of λ and are thus on an evolutionary indifference surface. These surfaces are linear functions of the x_i and are thus hyperplanes.

---

[24] The first-order difference equations describing the evolution of the population given by (3.3), (3.1) and (3.2) can be compactly represented in matrix notation as $N(t+1) = N(t)L$, where $N((t+1)$ is the population 3-vector at date $t+1$ and L is the 3x3 "Leslie matrix". The unique positive solution of the Euler-Lotka equation is the so-called "dominant eigenvalue" or "Frobenius root" of L. Frauenthal (1986) is a careful treatment of this result and its implications for demography.

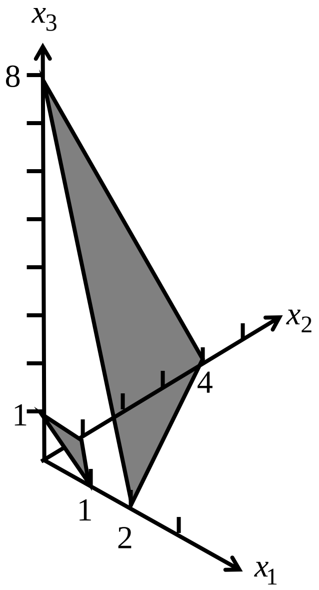

FIGURE 7. Hyperplanes on which λ is constant. This is Figure 1 from Robson and Samuelson (2007). [Copyright American Economic Association; reproduced with permission of the American Economic Review].

Since λ is constant, the Euler-Lotka equation implies the tradeoff between any two adjacent x_i as follows

$$0 = \frac{S_0 ... S_{i-1} dx_i}{\lambda^i} + \frac{S_0 ... S_i dx_{i+1}}{\lambda^{i+1}} \text{ so that } -\frac{dx_{i+1}}{dx_i} = \frac{\lambda}{S_i}.$$

3.2.1. *Preferences Over Intertemporal Consumption.* We now translate the above purely demographic account to come closer to the conventional intertemporal model of preferences.

The simplest way to connect resources to the demographic account is to suppose $x_i = f_i(c_i)$ gives age-i fertility as a function of age-i consumption c_i. The f_i are strictly increasing and concave.

For any consumption vector c, then,

$$1 = \frac{S_0 f_1(c_1)}{\lambda} + \frac{S_0 S_1 f_2(c_2)}{\lambda^2} + \frac{S_0 S_1 S_2 f_3(c_3)}{\lambda^3}, \tag{3.5}$$

where $\lambda > 0$ is constant on an indifference surface. $\lambda(c_1, c_2, c_3)$ is readily seen to be strictly increasing and can also be shown to be strictly quasi-concave in c. Hence the biological utility function λ generates indifference curves of the conventional type.

3.2.2. *Relationship to Conventional Model with Geometric Discounting.* Choosing the optimal consumption bundle (c_1, c_2, c_3) that maximizes the growth factor λ implies that the right hand side of the Euler-Lotka equation is maximized given the optimal λ. That is, suppose (c_1, c_2, c_3) is the optimal consumption profile and let λ^* be the associated maximal feasible growth rate. Then

$$1 = \frac{S_0 f_1(c_1)}{\lambda^*} + \frac{S_0 S_1 f_2(c_2)}{(\lambda^*)^2} + \frac{S_0 S_1 S_2 f_3(c_3)}{(\lambda^*)^3},$$

If (c_1', c_2', c_3') is a suboptimal profile generating growth factor $\lambda' < \lambda^*$ then

$$\frac{S_0 f_1(c_1')}{\lambda'} + \frac{S_0 S_1 f_2(c_2')}{(\lambda')^2} + \frac{S_0 S_1 S_2 f_3(c_3')}{(\lambda')^3} = 1 > \frac{S_0 f_1(c_1')}{\lambda^*} + \frac{S_0 S_1 f_2(c_2')}{(\lambda^*)^2} + \frac{S_0 S_1 S_2 f_3(c_3')}{(\lambda^*)^3}.$$

That is, the individual should maximize

$$\frac{S_0 f_1(c_1')}{\lambda^*} + \frac{S_0 S_1 f_2(c_2')}{(\lambda^*)^2} + \frac{S_0 S_1 S_2 f_3(c_3')}{(\lambda^*)^3}$$

given the *maximal* λ^*.

Mathematically, this is not very useful, because you need to know λ^* before you can use this observation to find the optimal c_1, c_2, c_3. But finding λ^* involves first finding the optimal c_1, c_2, c_3.

Nonetheless, this seems potentially relevant. That is, reasonable estimates of the modern pure rate of time preference often exceed estimates of mortality; certainly estimates of modern mortality away from the very beginning and end of life, but even mid-range estimates of hunter-gatherer mortality. So an extra source of impatience might help close the gap. We return to this issue below in Section 3.2.4.

3.2.3. *Time Preference in Hunter Gatherers.* An issue that arises with considering the evolutionary basis of time preference in hunter-gatherers is: Why would time preference have mattered if the technology did not permit borrowing or lending? That is, most food acquired by foragers must be consumed right away or it will be wasted.

A salient feature of hunter-gatherer societies, however, one that sets humans apart from chimpanzees, for example, is intergenerational transfer of resources. This is exhibited in Figures 8 and 9 which contrast male and female production and consumption of food. The enormous loan given by adults to the young is evident in the large initial excess of consumption over production followed by a counterbalancing excess of production over consumption. Chimpanzees, our closest living relatives, are close to autarky. This intergenerational transfer to the young is an enormous investment in skill acquisition, an investment that is amply rewarded by the adults' ultimate productivity. For this enormous investment to be worthwhile, it is also crucial that humans have low enough mortality. This is illustrated in Figure 10 which shows how combining human consumption with chimpanzee mortality would be catastrophic.[25]

These intergenerational transfers are not confined to immediate relatives. Sharing across all members of a band is most evident for meat obtained by collaborative hunting, usually by the males. One function of this sharing is risk-management, given that large animals are not inevitably successfully hunted. It may be that especially proficient hunters get the choicest parts, and it may be that larger families get somewhat less per child, but these effects arise as perturbations around largely equal division.[26] From an economics perspective, it is not at all obvious how such wide sharing could be sustained in equilibrium by selfish individuals.[27]

---

[25] See Robson and Kaplan (2006) for further discussion of the interaction between longevity and intelligence. The question of why aging occurs at all is addressed in Kaplan and Robson (2009).

[26] There is much less sharing of plants that are gathered.

[27] Another mostly hidden incentive might have been that proficient hunters were especially desirable mates.

But the brute fact is that it did happen. Hill and Hurtado (1996) provide a detailed account of food production and sharing among the Ache of Paraguay.[28]

In order to model these transfers, consider a hunter-gatherer society in steady state when all classes grow at the same factor λ^*. Consider the three adult age class example of the previous subsection. The steady state relative sizes of the three adult age classes are $\left(\frac{S_0}{\lambda^*}, \frac{S_0 S_1}{(\lambda^*)^2}, \frac{S_0 S_1 S_2}{(\lambda^*)^3}\right)$. Assume the population is large so that idiosyncratic risks of fertility or survival have essentially determinate effects.[29] That is, each age $a = 1, 2, 3$ class comes from an age 0 class a periods ago. Older age classes then stem from smaller age 0 classes as reflected in λ^*, and the size of the class is also deflated by survival. There is a social budget constraint at each time given by

$$\frac{S_0 c_1}{\lambda^*} + \frac{S_0 S_1 c_2}{(\lambda^*)^2} + \frac{S_0 S_1 S_2 c_3}{(\lambda^*)^3} = Y, \tag{3.6}$$

where

$$Y = \frac{S_0 y_1}{\lambda^*} + \frac{S_0 S_1 y_2}{(\lambda^*)^2} + \frac{S_0 S_1 S_2 y_3}{(\lambda^*)^3},$$

say. Although transfers over time are taken to be impossible, transfers between age cohorts at each point in time are unrestricted.

As we argued in Section 3.2.2, it is necessary that the optimal (c_1, c_2, c_3) (that maximize λ) should maximize the right hand side of Eq 3.5 subject to Eq 3.6 given the optimal λ^*. This particular case is simplified because the weights on the fertility function f in Eq 3.5 are the same as the weights on consumption c_a for $a = 1, 2, 3$, in Eq 3.6. In the case that $f_i = f$, for example, this implies that the c_1, c_2, c_3 that maximize the right hand side of Eq 3.5 subject to Eq 3.6 are equal.

The evolutionarily optimal allocation of resources is then simply described. Pool all the resources contributed by all the living representatives of each age cohort and reallocate them equally.

---

[28] See also Robson and Kaplan (2003) for discussion of this issue.

[29] Finite group sizes would make the problem substantially more awkward. Hunter gatherer groups were typically around 30, which would make the resulting stochasticity of transfers small but still possibly important.

It is as if the individual were maximizing a conventional lifetime utility function subject to a conventional lifetime budget constraint, one that allows borrowing and saving, and then choosing a substantially different consumption pattern from the production pattern. These intertemporal preferences entail a pure rate of time preference that arises from survival and from the population growth rate, and are as derived earlier in Section 3.2.2.[30]

In this sense, hunter-gatherer economics generates two key features of conventional intertemporal preferences that otherwise seem arbitrary—additive separability and a constant discount factor, once mortality is netted out. These key features trace back to the application of the Euler-Lotka equation.

These preferences might have been explicitly expressed among hunter-gatherers if individuals were able to make small individual changes in consumption levels at various ages. There were suitable privately owned assets—bows and arrows, for example—even among the Ache, who were among the more egalitarian of hunter gatherer societies. These assets would have permitted modest intertemporal transfers that parallel those available in modern societies.[31]

---

[30] The results sketched here can be substantially generalized. Robson and Samuelson (2022, Section 5) consider endogenous fertility and mortality when intergenerational transfers are freely permitted, but all social intertemporal transfers are ruled out. Mortality can be reduced by resources devoted to that purpose. Fertility is assumed to be subject to a threshold effect that makes zero fertility optimal at very young and old ages. It is then optimal to invest resources in people who have no future fertility as long as they retain future output, as in the "grandmother effect". Each individual trades off resources available at various ages with the same pure rate of time preference as found here. These results are part of an examination of the effect of occasional aggregate shocks, but this involves a complete treatment of the case with no such shocks.

[31] Manvir Singh uses these to cast a skeptical light on the standard view in anthropology of hunter-gatherers as egalitarian—https://aeon.co/essays/the-idea-of-primitive-communism-is-as-seductive-as-it-is-wrong. Data from the Ache on massive intergenerational transfers that nevertheless reinforce the standard view are provided by Hill and Hurtado (1996).

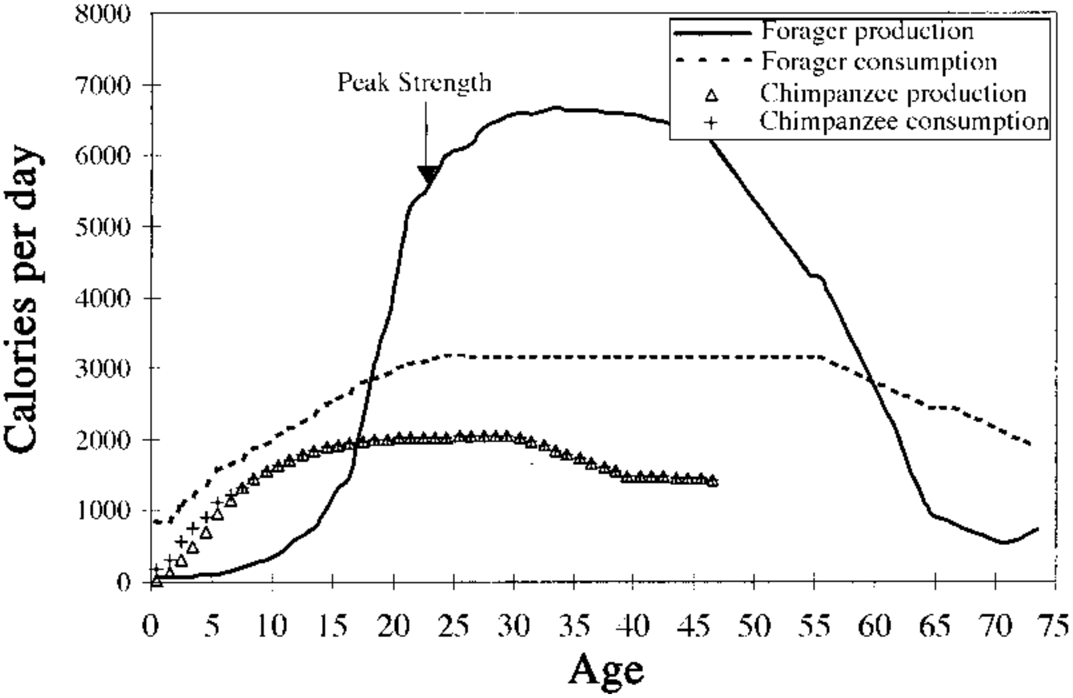

FIGURE 8. The figure describes male production and consumption of food, contrasting humans and chimpanzees. That male production peaks noticeably after peak strength speaks to the role of skill and intelligence in hunting. The human data are for the Ache, Hiwi, and Hadza, who remained as hunter-gatherers into the 20th century; the chimpanzees are in the wild. The relatively minor transfer from adult chimpanzees to the young is mainly mother's milk. This reproduces Figure 2A from Robson and Kaplan (2003), which gives precise references. [Copyright American Economic Association; reproduced with permission of the American Economic Review].

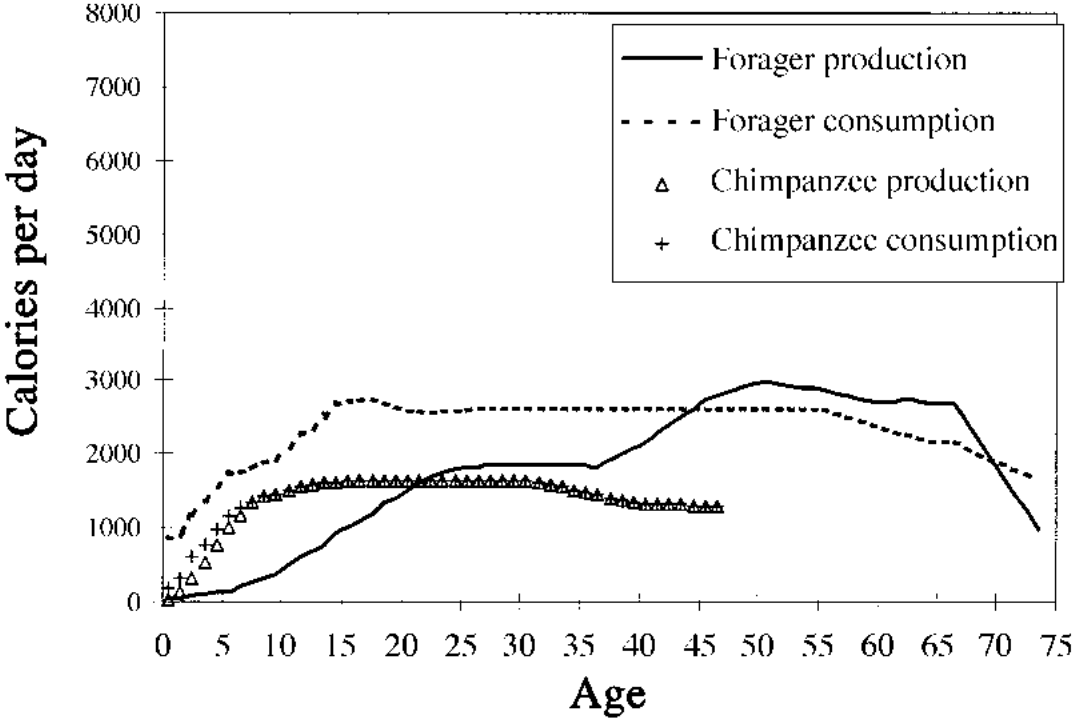

FIGURE 9. The figure describes female production and consumption of food, contrasting humans and chimpanzees. Compare with the male schedules. The significant contribution by females in rearing children is missing from this picture. This reproduces Figure 2B from Robson and Kaplan (2003). [Copyright American Economic Association; reproduced with permission of the American Economic Review].

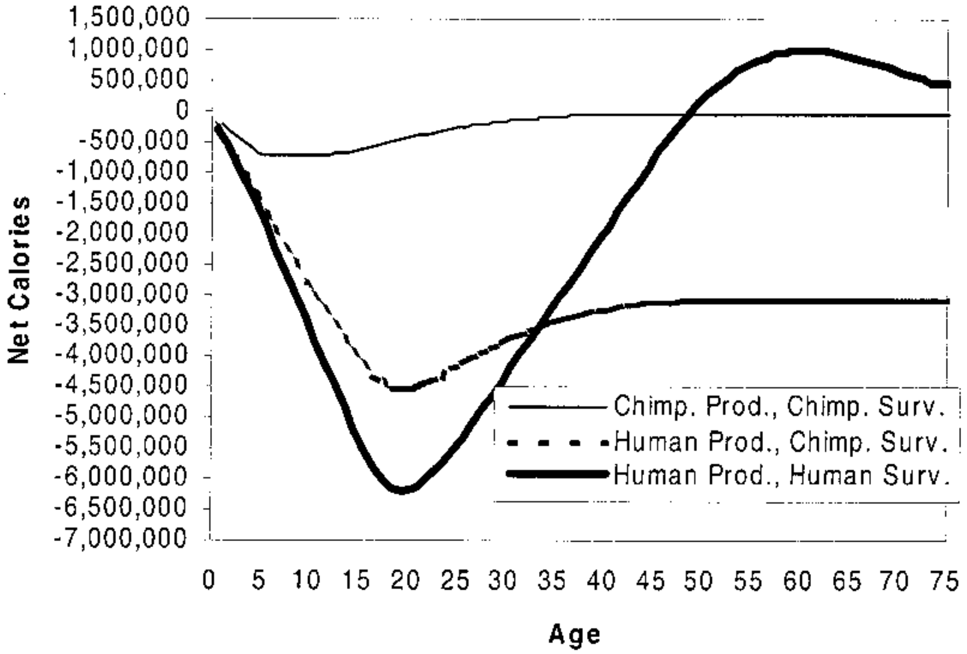

FIGURE 10. This describes cumulated expected net production of food by age, contrasting humans and chimpanzees. At each age, the difference between production and consumption, weighted by the probability of the individual still being alive, is added to the running total. which was initialized at zero. This calculation corresponds to finding the net social surplus in a society with free transfers across ages in a steady state with no growth. Human adults are able to cover the expected deficits of the young in the steady-state, as is revealed by this cumulated excess converging to a small positive level, which would actually permit moderate growth. This requires low mortality. To dramatize this, suppose human net production is combined with chimpanzee mortality. In this hypothetical case, the cumulated expected excess remains well below zero. In the steady state, such a deficit could only be accommodated with a catastrophic population decline. This reproduces Figure 4 from Robson and Kaplan (2003). [Copyright American Economic Association; reproduced with permission of the American Economic Review].

3.2.4. *The Effect of Population Growth on Time Preference.* The present approach provides a basis for the usual simple model of time preferences. There is a component of the pure rate of time preference, involving population growth, which entails a *constant* rate of discount of the future, justifying what seems an arbitrary assumption made in the simplest model of intertemporal preferences in economics

The influence of mortality on time preference is well-appreciated in economics (see Fisher (1930), for example). Although the usual approach does not emphasize age-varying time preference, a possibly age-related effect of mortality is then as expected. What the present derivation points out is that there is an age-constant influence of the rate of population growth on time preference.

However, the average λ over the 1.8 million years of our evolutionary history must be very close to 1, as a matter of arithmetic necessity. It follows then that the pure rate of time preference, $\ln \lambda - \ln S = \ln \lambda + \delta$, is simply mortality δ. So population growth would seem unable to contribute to an explanation of the pure rate of time preference. The puzzle of the remaining gap between the biological prediction of the rate of time preference and the rate that is observed is the motivation for Robson and Samuelson (2008), which we take up next.

This closing of the gap invokes aggregate risk concerning survival. The differential effect of idiosyncratic and aggregate risk is examined in greater detail in the subsequent Section 4.

Reconsider the three age population model of Sections 3.2.3 and 3.2.4. There is a population structure at t given by $(N_1(t), N_2(t), N_3(t))$. All survival probabilities are subject to an *aggregate* shock[32] so that, with three adult ages—

$$N_1(t+1) = \tilde{S}\left\{N_1(t)x_1 + N_2(t)x_2 + N_3(t)x_3\right\}$$

$$N_2(t+1) = \tilde{S}N_1(t).$$

$$N_3(t+1) = \tilde{S}N_2(t)$$

where $\tilde{S}$, the survival factor, is i.i.d.

---

[32] It is easy to allow arbitrary idiosyncratic risk as well.

Consider also the evolution of the population stripped of all *interim* mortality

$$M_1(t+1) = M_1(t)x_1 + M_2(t)x_2 + M_3(t)x_3$$
$$M_2(t+1) = M_1(t).$$
$$M_3(t+1) = M_2(t)$$

Suppose that the stripped population has steady state growth factor λ. The stripped population is in steady state growth if and only if the original population is in steady state growth. This is because the imposition of a uniform survival factor does not shift the age structure of the population. Think of the original population as first growing without mortality followed by the imposition of mortality. Then, in the first period, the growth factor is just

$$\lambda \tilde{S},$$

so that the growth of the original population is the stripped growth factor multiplied by the random survival factor.

Consider the exponential rate at which iid products of the survival factor $\tilde{S}$ shrink. Suppose, to take a simple case, that $\tilde{S} = S_1$ with probability p and $\tilde{S} = S_2$ with probability $1 - p$ independently over time. Hence the overall survival factor after T periods is $\tilde{S}(T) = S_1^{T_1} S_2^{T-T_1}$, where T_1 is the number of times $\tilde{S} = S_1$. Hence $\frac{1}{T} \ln \tilde{S}(T) = \frac{T_1}{T} \ln S_1 + \frac{T-T_1}{T} \ln S_2 \to p \ln S_1 + (1-p) \ln S_2 = E \ln \tilde{S}$, as $T \to \infty$. This result readily generalizes to any number of outcomes and to continuous distributions.

The results of Robson and Samuelson (2008) can then be given in a highly simplified form as follows. The growth rate of the factor $\lambda \tilde{S}$ is now

$$\Lambda = E \ln \tilde{S} \lambda(x),$$

so that

$$-\frac{dx_{i+1}}{dx_i} = \frac{\frac{d\Lambda}{dx_i}}{\frac{d\Lambda}{dx_{i+1}}} = \frac{\frac{d\lambda}{dx_i}}{\frac{d\lambda}{dx_{i+1}}} = \lambda.$$

Thus the discount rate is the growth rate that arises with no mortality. (This is equivalent to the claim that the discount rate is the sum of the actual growth rate plus the mortality rate.)

The current approach with aggregate survival can accommodate realistic pure rates of time preference if it is first possible to construct a comparable growth rate in a hypothetical case where mortality is set to zero. Secondly, it should be possible to incorporate plausible aggregate mortality capable of bringing the realized growth rate back to zero.

We then need an idea of how fast a population would grow in the absence of any mortality. Consider the following example. Suppose that individuals start reproducing at age 15 and stop at age 45, with no risk of death before age 45. Hill and Hurtado [1996, Chapter 8, especially Table 8.3] study of the Ache suggests a prime-age birth probability of 0.15 per year. (This cuts the birth probabilities reported there in half. The 0.15 then represents the probability of a female birth, providing a valid comparison with the model of asexual reproduction.)

Since individuals start reproducing at age 15 and stop at age 45, and the probability of giving birth in a given year is 0.15, the dominant eigenvalue solves $1 = \sum_{\tau=15}^{45} \frac{(.15)}{\lambda^t}$, so $\lambda = 1.05675$ and $\ln \lambda = 0.055$. If the growth rate is zero,

$$0 \;=\; \Lambda \;=\; \ln \bar{\lambda} + E \ln \tilde{S} \;=\; 0.055 + E \ln \tilde{S},$$

Suppose further that, with probability $1 - p$, the death rate is about two percent, reflecting the idiosyncratic risk of dying from accidents and disease. With remaining probability p, an aggregate catastrophe with survival rate of $S^{\dagger}$ appears. Then $E \ln \tilde{S} \;=\; p \ln S^{\dagger} + (1 - p)(-0.02) \;=\; -0.055$. Thus if $S^{\dagger} \;=\; 0.50$, then $p \;=\; 0.05$ works for example. Alternatively, if $p = 0.01$ then $S^{\dagger} = 0.03$.

It could easily be that events with probabilities of only 1% would entirely escape notice in the geological record of humans. This would mean that human populations could generate a pure rate of time preference as high as 5.5% despite the necessity of a long run rate of growth near zero and an apparent mortality rate of perhaps only 2%. Even if the possibility of such events were recognized, it would add only about 1% to the raw probability of an individual dying each year. (The

1% is an upper bound, arising from estimating that the catastrophe kills everyone.) The aggregate nature of the catastrophe leads to an exaggerated effect of about 3.5% on the long run growth rate. With no recognition of this effect, that is, there would still be 2.5% unaccounted for. The lower the probability of the catastrophe, the greater the exaggeration, assuming the catastrophe remains severe enough to keep the overall long run growth rate at zero. It is essentially this exaggeration that we explore further in Section 4, where we examine the differential evaluation of aggregate and idiosyncratic risk in preferences.

The catastrophe described in the preceding paragraph sounds suspiciously severe, but there is amazing evidence, drawn from the DNA of modern humans, that, somewhat less than a million years ago, an event occurred in which only about one percent of the population of human ancestors survived, reducing the population from about 100,000 to about 1000, with the population remaining at the latter level for about 100,000 years—see Hu, Hao, Du, Di Vincenzo, Manzi, Cui, Fu, Pan, and Li (2023). Such an event would have made our ancestral line teeter on the brink of extinction. The event seems to have driven a marked reduction in human genetic diversity and the creation of a new human species. The period apparently experienced an unfavorable climate shift and coincides with a substantial gap in the fossil record.[33]

That is, this argument is readily capable of generating a pure rate of time preference of about 5% which covers most of the non-experimental empirical observations on the pure rate of time preference. It covers, for example, the moderate values used by Nordhaus (1994) in his analysis of global warming. Such moderate values are enough to tilt global warming away from being a dire problem to being more tractable. Stern (2006) derives the dire consequences of taking the pure rate of time preference to be near zero. The issue for Stern, however, is not so much to estimate the pure private rate of time preference but to argue, on philosophical grounds, that the public rate should be near zero.

---

[33] The prolonged nature of this episode illustrates the need to investigate environmental shocks that are *not* i.i.d.

It is true that even carefully done experiments may produce substantially higher estimates of time preference than 5%—see Harrison, Lau, and Rutström (2007) , for example. The rates of time preference derived experimentally, however, often do not seem credible outside the laboratory. Oprea (2024) suggests that experimental measurement of risk aversion might actually be measurement of aggregation frictions. Perhaps the experimental measurement of time preference is also subject to an alternative interpretation in terms of cognitive or perceptual frictions.

3.2.5. *Plasticity of Time Preference.* The argument concerning adaptation of attitudes to risk in Section 3.1.1 can be extended to time preference. That is, if a shift in the environment leads to a change in the demographic variables, and this shift is signaled, individuals should make different intertemporal trade-offs, reflecting the pure rate of time preference suited to each environment. Modeling time preference entails a nontrivial population structure and this makes the argument less straightforward now, however, because the process of converging to a new steady state—one in which the pure rate of time preference applies—will take time. Nevertheless, at least if the shift in the environment is rare, the advantage of adaptation is clear.

To see this, consider the simplest possible demography, with just two adult ages $a = 1, 2$. Suppose survival from age 0 to 1, or from age 1 to 2 is certain, for further simplicity. There are two states, $s = 1, 2$, and two actions, $a = 1, 2$. Fertility at either adult age if action a is taken in state s is $u(a, s) > 0$, which is the expectation of idiosyncratic risk conditional on the aggregate state s and action a. Suppose that action a is optimal in state a for $a = 1, 2$, so that $u(1, 1) > u(2, 1)$ and $u(2, 2) > u(1, 2)$. The long run growth factor from action a in state s is then λ_{as} which is the unique positive solution for λ of

$$1 = \frac{u(a,s)}{\lambda} + \frac{u(a,s)}{\lambda^2}.$$

That is, using the quadratic formula,

$$\lambda_{as} = \frac{u(a,s) + \sqrt{u(a,s)^2 + 4u(a,s)}}{2}.$$

Of course, it follows that $\lambda(1, 1) > \lambda(2, 1)$ and $\lambda(2, 2) > \lambda(1, 2)$. Suppose without much loss of generality that $u(1, 1) > u(2, 2)$, so that taking the appropriate action in state 1 outperforms taking the appropriate action in state 2 and hence the pure rate of time preference is higher.

Suppose now the state is driven by a simple Markov process in continuous time with a continuous time probability rate $\pi > 0$ of making a transition either from state 1 to 2 or vice versa. It follows that the limiting probability of either state is $1/2$. Consider now the further limit as $\pi \to 0$.[34] Since transitions are then rare, the population has time to settle down to grow at rate $\ln \lambda_{as}$ if action a is taken in state s. If the state cannot be seen by individuals then the limiting overall growth rate from choosing action $a = 1, 2$ is

$$\frac{\ln \lambda_{a1} + \ln \lambda_{a2}}{2}.$$[35]

If the state can be observed, the individual can and should choose the matching action, generating overall limiting growth rate

$$\frac{\ln \lambda_{11} + \ln \lambda_{22}}{2}$$

which is clearly higher than either limiting growth rate attainable when the state is not observed.[36]

Hence, it is advantageous that time preference—which is reflected in the factor λ here—to be plastic and adapt rapidly to the environment. The genetic mechanism that enables such rapid choice of the matching action might itself be subject to slow natural selection, and would require repeated exposure to the various environments. That is, natural selection does not amount to a straitjacket that limits the range of observable time preferences.

---

[34] This assumption goes beyond those to obtain the plasticity of risk aversion in Section 3.1.1. There was no need there to allow the population structure to settle down.

[35] The mathematics needed to show this claim rigorously are as in Robson and Samuelson (2019).

[36] Again, the treatment of the aggregate states $1, 2$ here anticipates the treatment of aggregate gambles in Section 4. The options here involve intertemporal tradeoffs derived from idiosyncratic risk *conditional* on each aggregate state.

3.2.6. *Ancestral Environment and Modern Time Preference.* Estimates of the pure rate of time preference vary substantially by region and country. Falk, Becker, Dohmen, Enke, Huffman, and Sunde (2018) (FBDEHS) report and analyze an extensive survey of economic attributes including time preference and risk preference from 80,000 people in 76 countries. They relate these attributes to contemporaneous variables such as geography, culture, language, and religion. Becker, Enke, and Falk (2020) look back in time, and show that populations that diverged genetically due to earlier migration exhibit more pronounced differences in economic attributes.[37]

What does the current evolutionary argument imply about how the ancestral environment would shape the genes that influence modern time preference? In particular, what ancestral demography might indirectly generate high modern pure rates of time preference?

To address this, consider the three-adult-age-class model under the additional simplifications that all adult fertility levels are identical at f and that all survival probabilities are also identical at $S \in (0,1)$.[38] The Euler-Lotka equation is then

$$1 = \frac{Sf}{\lambda} + \frac{S^2 f}{\lambda^2} + \frac{S^3 f}{\lambda^3}.$$

The factor reflecting the pure rate of time preference, which discounts fertility, is then $\frac{\lambda}{S}$. If there were data on these parameters for the ancestral environment, it would be possible to verify whether high estimates of $\frac{\lambda}{S}$ were associated with high modern pure rates of time preference.

Even if estimates of ancestral population growth are missing, a rough estimate of population growth follows as matter of arithmetic necessity. That is, over any long period of time, thousands of years, say, the ancestral population could only have grown at a low rate, so that $\lambda = 1$, approximately.[39] If $\lambda = 1$, a higher pure rate of time preference would then arise from lower survival S or, equivalently, from

---

[37] Galor and Özak (2016) also look back in time to test their hypothesis that agriculture, in particular, shaped the pure rate of time preference.

[38] These assumptions can be considerably relaxed. The simplicity here is meant to capture the central effect of mortality, for example, on time preference, as might be investigated empirically.

[39] A complete account of fertility and mortality would incorporate population as an argument so that $\lambda \to 1$ as the population approaches carrying capacity.

higher mortality. This would necessitate higher fertility f.[40] Hence an ancestral regime that would lead eventually to high modern pure rates of time discount might have had high fertility and high mortality.

In the data presented by FBDEHS (2018, Figure 1, for example), Sweden is one of most patient countries that have been settled for a long time. Did Sweden experience the low mortality and low fertility that would induce high patience? Did Egypt, which is an old country that features as impatient, experience high mortality and high fertility?

4. ATTITUDES TO AGGREGATE RISK

In the preceding section, we dug into the evolutionary foundations of attitudes to idiosyncratic risk and time preference, applying a demographic model of literal natural selection. We anticipated there the asymmetry between the effect of idiosyncratic risk and that of aggregate risk. We now focus further on this asymmetry. In particular, using a demographic model, we examine the prediction that individuals have different attitudes to aggregate risk—where a single giant coin is flipped to decide an outcome for everyone—versus idiosyncratic risk—where individuals flip their own personal coins.

We thus derive preferences that are not purely selfish—to evaluate an agent's preferences, it is not enough to know the distribution of that agent's offspring alone. The correlation with other agents' offspring outcomes is also important. This interdependence arises despite a resolutely selfish gene perspective. Success is inevitably measured in an evolutionary sense by a large number of adherents, but a large number of adherents passes through the filter composed of a sequence of aggregate states to an extent that cannot be captured by selfish preferences. This is an intriguing challenge to providing an evolutionary basis for purely independent preferences in all circumstances.

The results of Robson (1996), which frames the question in discrete time, can be illustrated with a simple example. Individuals of Type I have either $f_1 > 0$ or $f_2 > 0$ expected offspring each with probability $1/2$. If all this risk is idiosyncratic,

---

[40] Indeed, in this case, $f = \frac{1}{S+S^2+S^3}$.

and there is a large number of individuals, the population grows according to the mean. That is, if $x(T)$ is the number of Type I at period $T = 1, 2, ...$, with a population of size 1 initially, then

$$x(T) = \left(\frac{f_1 + f_2}{2}\right)^T, \text{ so } \frac{\ln x(T)}{T} \to \ln\left(\frac{f_1 + f_2}{2}\right) \text{ as } T \to \infty,$$

with probability 1, in the light of the strong law of large numbers.

Suppose, on the other hand, individuals of Type II experience the same distribution of outcomes, but as aggregate risk. That is, if $y(T)$ is the number of Type II's at date $T = 1, 2, ...$, then $y(T) = f_1^{n(T)} f_2^{T-n(T)}$, where $n(T)$ is the random number of heads in T flips of a fair coin. Hence

$$\frac{\ln y(T)}{T} = \frac{n(T)}{T} \ln f_1 + \frac{T - n(T)}{T} \ln f_2 \to \frac{\ln f_1 + \ln f_2}{2}, \text{ as } T \to \infty,$$

with probability 1.

Since $\frac{\ln f_1 + \ln f_2}{2} < \ln\left(\frac{f_1+f_2}{2}\right)$, by concavity of the log function, individuals exposed to aggregate risk experience lower growth rates than those exposed to precisely comparable idiosyncratic risk. That is, individuals should be more averse to aggregate risk than they are to idiosyncratic risk.

Robatto and Szentes (2013) cast the question in continuous time instead and obtain contrasting results. Consider now Type III, say, that produces offspring either at a rate of $r_1 > 0$, which applies in State 1, or $r_2 > 0$, which applies in State 2. The state is governed by a continuous time Markov process with a probability of $\pi \in (0, 1)$ per unit time of making a transition in either direction. The limiting probability of being in either state is then $1/2$. If $z(T)$ is the number of Type III at time T, and there is a population of size 1 initially, then

$$z(T) = e^{r_1 T_1} e^{r_2 (T - T_1)}$$

where $T_1 \in [0, T]$ is the length of time that State 1 is in effect. It follows that

$$\frac{\ln z(T)}{T} = r_1 \frac{T_1}{T} + r_2 \frac{T - T_1}{T} \to \frac{r_1 + r_2}{2} \text{ as } T \to \infty,$$

with probability one.

Hence individuals exposed to this aggregate risk experience a random growth rate, but one generating the same limiting growth rate as the mean of the random growth rates. Individuals then should be indifferent to the choice between the aggregate gamble over growth rates and a deterministic growth rate that is the mean of the random growth rates.

The demographic structure of the Robatto and Szentes model is stark. All individuals regardless of age experience the same underlying fertility and mortality rates. Robson and Samuelson (2019) show that this stark structure lies behind the risk-neutrality obtained. Once age affects fertility or mortality rates, there is generally an evolutionary distinction between the effects of aggregate and idiosyncratic risk. Usually, but not always, idiosyncratic risk is preferred to comparable aggregate risk. If there is an age gap before fertility begins, as seems only realistic, this gap arising from age dependent demography has effects roughly comparable to the gap formally created in the discrete time model. In the simplest case, individuals reproduce as a spike at one particular age, in which case the continuous time model generates essentially the same results as the discrete model.[41]

In the final analysis, the issue cannot be resolved on a purely theoretic basis. Rather, it is: which of the models best approximates reality? That is, do people exhibit a preference for idiosyncratic risk over comparable aggregate risk? This is the question we look at in the present section.

In environments where agents freely choose reproductive gambles, any effect on the population growth rate of aggregate risk will be reflected in individual attitudes to aggregate risk. Curry (2001), for example, considers a discrete time model (as in Robson (1996)), but where each individual chooses a gamble that is a best response to the distribution of gambles of the others in the population. The evolutionarily dominant type is shown to be one that maximizes relative numbers of own offspring, that is, own offspring divided by offspring per individual for the remainder of the population. Evolutionarily dominant preferences are interdependent, and importantly, individuals are averse to correlation in own offspring and average offspring of the rest of the population.[42]

---

[41] Robson and Samuelson (2024) includes an extensive discussion of this debate.

[42] Heller and Nehama (2022) provide another utility based implementation of the evolutionary optimum in a discrete time set-up. In their dual-utility approach individuals use one utility function to

Correlation aversion here is in the sense of Epstein and Tanny (1980), as follows. Consider joint distributions f and g over a pair of variables (x, y). Define f as inducing more correlation than g if f and g imply the same marginal distributions over x and y, but under f it is more likely to observe pairs (x, y) where x and y are both low, or both high.[43] An individual is then averse to correlation in x and y if she prefers the distribution g to f whenever f induces more correlation than g. Epstein and Tanny (1980) shows that an individual with utility over pairs, (x, y), is averse to correlation in x and y if, and only if, her marginal utility of x decreases in y. Aversion to correlation in offspring then follows immediately in Curry (2001), since there the marginal utility of own offspring decreases in the average number of offspring of the others in the population.

Oprea and Robalino (2024) use the evolutionary theory to generate testable hypotheses regarding attitudes to correlation in consumption. As in Robson (1996), consider a biological population where consumption is mapped to offspring as follows. If an individual consumes x, then she has j offspring with probability $p(j, x)$. Assume realized own offspring is independent across individuals and independent of the state, conditional on own consumption. Assume more consumption unambiguously improves the distribution of offspring, in the sense of first-order stochastic dominance.

The following is then proved by Oprea and Robalino (2024). If aggregate risk in offspring is evolutionarily disadvantageous (as in the model of Robson (1996)), then this will be reflected in individuals having an aversion to correlation of own

---

assess idiosyncratic gambles over offspring, and another more risk averse utility function to assess gambles entailing aggregate risk, where there is heterogeneity in preferences over correlated gambles.

[43] More formally, f induces more correlation in x and y than g if they imply the same marginal distributions, and for all $\alpha, \alpha' \geq 0$,

$$f(\{x \geq \alpha, y \geq \alpha'\}) \geq g(\{x \geq \alpha, y \geq \alpha'\}),$$

and

$$f(\{x \leq \alpha, y \leq \alpha'\}) \geq g(\{x \leq \alpha, y \leq \alpha'\}).$$

These inequalities, holding for all α, and α', mean that under f it is more likely to observe pairs (x, y) where x and y are both high, or both low, hence f *induces more correlation* than g.

and others' consumption. On the other hand, if aggregate risk in offspring is evolutionarily neutral (as in the model of Robatto and Szentes (2013)), then this will be reflected in individuals being indifferent to correlation of own and others' consumption. Finally, if aggregate risk in offspring is evolutionarily advantageous, then this will be reflected in individuals having an affinity for correlation of own and others' consumption.

In a large population the effect of correlation in own consumption and the consumption of a small number of others is negligible. Hence the results here describe preferences for correlations between own consumption and that of a large number of others. Whatever attitude is evolutionarily appropriate might, however, be reflected in corresponding attitudes in the small. For example, if aversion to correlation is evolutionarily advantageous, a heuristic might arise to be averse to correlation in own consumption and the consumption of a small group of others. Oprea and Robalino (2024) test correlation attitudes in small groups.

4.1. **Experimental Evidence.** Oprea and Robalino (2024) experimentally test the implications of the evolutionary theory. Their test is based on the bomb risk elicitation task by Crosetto and Filippin (2013). The bomb task offers an intuitive method for measuring risk attitudes. A subject decides how many boxes to stack, $b \in \{0, 1, 2, \ldots, 50\}$. A laser then fires, randomly and uniformly, from a vertical position $s \in \{0, 1, 2, \ldots, 50\}$. If the laser misses the subject's stack, then she earns some amount m for each box stacked. If the laser strikes the subjects' stack, then she earns zero. Subjects that are more risk averse (loving) will stack fewer (more) boxes. A risk neutral subject will choose $b = 25$. Risk averse subjects will stack $b < 25$ boxes, and risk loving subjects will choose $b > 25$. The test allows between subject comparisons. If subject A is more risk averse than B, in the sense of everywhere having a higher Arrow-Pratt coefficient of risk aversion, then A stacks fewer boxes than B.

In the ID (idiosyncratic) task each subject makes a choice in the bomb task, alone, in order to measure risk preferences in an uncorrelated environment. Each subject also faces one of three social bomb tasks. In the LEAD social task the subject chooses a stack height for herself, and this choice determines the stack height of

30 other subjects. All 31 subjects then face the same laser. Hence, payoffs are perfectly correlated—all the subjects involved in the task earn zero, or the payment implied by the decision maker's choice.

Importantly, for each given box choice, the distribution of the subject's payoff is the same in ID and LEAD. An indifference to correlated payoffs (as in conventional expected utility theory) implies the same choice in the two tasks. A greater aversion to aggregate risk than to commensurate idiosyncratic risk implies a less risky choice in the LEAD than in ID.[44]

Results from ID fall within the normal range of measured risk aversion in the prior literature. The mean subject choses $b = 19.5$, which corresponds to a CRRA parameter of 0.738, given the utility specification $u(x) = x^r$. In LEAD, however, the implied CRRA parameter decreases by 0.0118. The drop is statistically significant, and compares in magnitude to the difference in risk aversion recorded between men and women. Hence, the resulting decrease in risk taken in LEAD relative to ID is consistent with greater aversion to aggregate risk than to commensurate idiosyncratic risk.

The difference in risk taking observed between LEAD and ID, however, might also be explained by some other type of interdependent preferences, rather than aversion to aggregate risk. For example, subjects might have felt uncomfortable taking risks for others. They might have then tailored their choices in LEAD to accommodate what they believed were the risk preferences of the other subjects. If this were the case, then the reduced risk taking observed in LEAD relative to ID, would not be due to correlation aversion. Two other treatments are carried out in order to address such a concern. In each of these, a subject only chooses for herself, but her payoffs are correlated with the payoffs of other subjects (positive correlation in POS and negative correlation in NEG). We discuss these in detail below.

In the POS (positive correlation) social task, the decision maker chooses her own stack height, but is shown the box height choices of 30 other subjects. All the subjects face the same bomb, as in LEAD. Payoffs in POS are thus positively correlated

---

[44] The proof is given by Oprea and Robalino (2024).

(see Figure 11). In POS, however, in contrast to LEAD, the subjects' choice has no implications for the payoffs of the others.

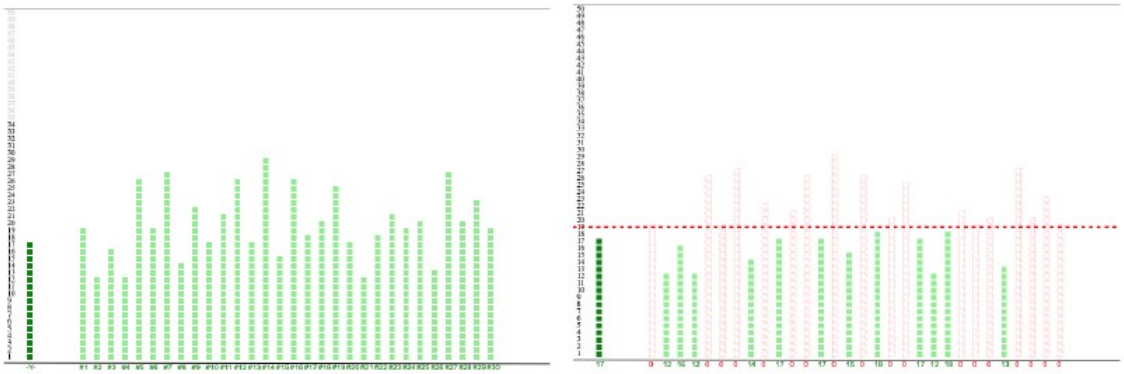

(e) POS Screenshot (During Task)

(f) POS Screenshot (After Task)

FIGURE 11. On the left is the view of a subject in the POS task. The dark green stack to the far left is the subject's stack. The other stacks are stacks of the other subjects. The subject's choice does not affect the payoffs of the others. On the right is the subject's view after the laser has been fired. Subjects with stacks struck by the laser (faded stacks) get zero. The others are paid according to their chosen box heights.

Again, if a subject is indifferent to correlation in own payoff and the payoffs of the others, as in the conventional theory, then her choices in POS and ID will be the same. If the subject is averse to correlation in own and others' payoffs, then she will make a less risky choice in POS than in her ID task (this is proved in Oprea and Robalino (2024)). On the other hand, an affinity for such correlation implies a more risky choice in POS than in ID. Results from the POS treatment are in line with those from LEAD, suggesting an aversion to correlated payoffs.

A third treatment, NEG, controls for possible social imitation in POS. In NEG, the set-up is as in POS (Figure 11), except that the decision maker's stack is hung from the ceiling. The situation is identical to POS except that the subject's payoffs are now negatively correlated with the payoffs of the 30 other subjects. Again, the conventional theory predicts the same choice in NEG and ID. An aversion to (affinity for) correlated payoffs implies a more (less) risky choice in NEG than in ID. Subjects do make slightly more risky choices in NEG than in ID, in line with aversion to correlation, though the difference is not statistically significant. In any case, negative correlation eliminates the increased risk aversion observed when payoffs are positively correlated, suggesting it is positively correlated payoffs that produce more risk aversion in POS than in ID.

Results for all treatments are summarized in Figures 12 and 13. They contradict conventional expected utility theory, where individuals care only about the marginal distribution of own payoffs. In particular, positively correlated payoffs induce more risk averse choices, and the effect is shut down when payoffs are negatively correlated. Hence, more risk averse choices in LEAD and POS appear to be driven by positively correlated payoffs. This is consistent with a greater aversion to aggregate risk than to commensurate idiosyncratic risks, and with a corresponding aversion to correlation in own and others' consumption, in line with the model from Robson (1996).

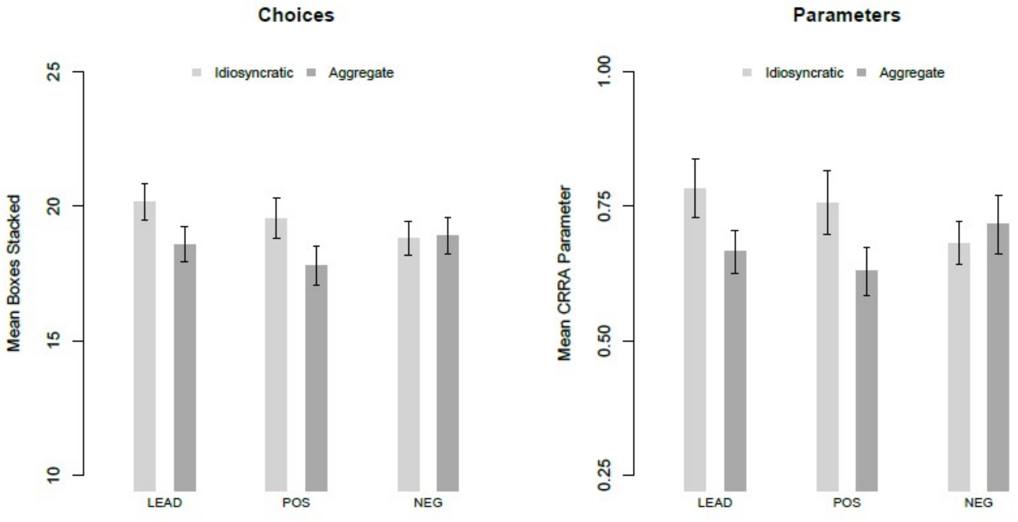

FIGURE 12. On the left are box heights observed in each ID task and in the corresponding LEAD/POS/NEG task. On the right is the implied CRRA parameter in each ID task and the corresponding LEAD/POS/NEG task.

Figure 12 shows the average ID choice is riskier than the average choices in the LEAD and POS task. Risk taking is higher in NEG than in ID, but the difference is not statistically significant. Importantly, within-subject comparisons are consistent with correlation aversion, which is implied by the biological theory that aggregate risk in offspring is disadvantageous (Robson (1996)). Figure 13 shows the difference, averaged across subjects, between an individual's choice in ID and her corresponding choice in the social task (LEAD, POS, or NEG). The mean subject makes a riskier choice in ID than in LEAD and POS. The difference is significant. Again, the mean subject makes a riskier choice in NEG than in ID, though the difference is not statistically significant.

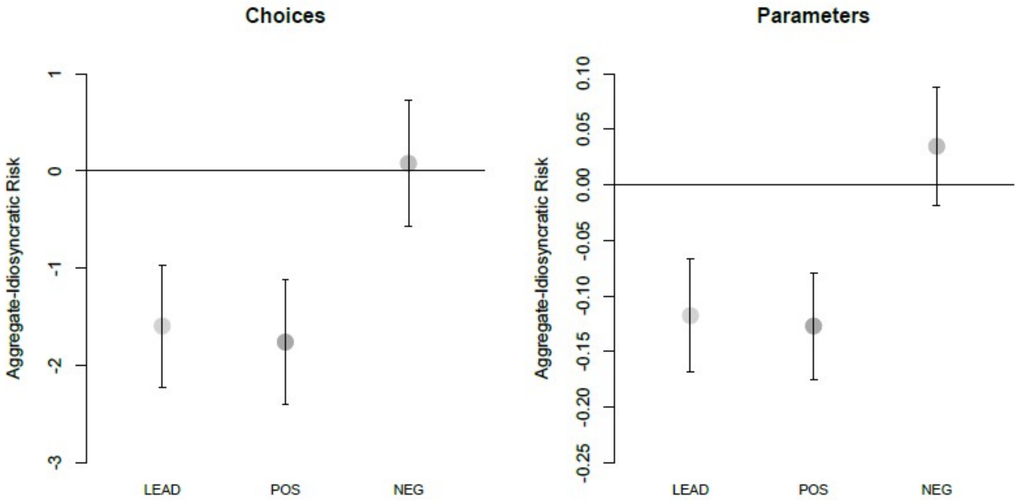

FIGURE 13. On the left is the within-subject difference in box height between the ID task, and LEAD/POS/NEG. On the right is the inter-subject difference in implied CRRA parameter between ID and LEAD/POS/NEG.

Age-structured models predict interesting patterns of attitudes to correlation in offspring and mortality across different age cohorts (for example, see Robson and Samuelson (2019)), suggesting further experimental work is needed.

This section has sketched how a striking theoretical prediction of the current evolutionary approach—that individuals are more averse to aggregate risk than to idiosyncratic risk—finds empirical validation. That this prediction came from left field—would not otherwise have been readily foreseen—reinforces the scientific merit of the approach.

5. THEORY OF MIND

We have so far discussed the evolution of utility, and some of its key properties, thereby providing a biological basis for economic decision-making in non-strategic situations. Humans are highly social, however. Many important choices involve interaction with others. In conventional game theory we are assumed to be fully sophisticated. Individuals in a game attribute motives and intentions to others, and they fully understand the implications of these mental states. This is, moreover, understood by all the players in the game—that is, player 1 is intentional, knows that player 2 is intentional, knows that player 2 knows that player 1 is intentional, and so on. There are alternatives to this fully strategic conception of social interaction. An agent in a game need not be aware at all of others' mental states, but instead might learn the appropriate choice through reinforcement learning, for example. There is good evidence, however, that humans are strategic, at least to some extent—that they reason about the motives and beliefs of others. We now explore the evolution of this strategic sophistication.

Theory of mind (ToM) refers to the ability to impute mental states to oneself and to others. An individual with ToM can conceive of others as intentional beings that are motivated by internal states. She can infer these states from observed behavior, and can use these inferences to predict behavior in the future. The term "theory of mind" was first used by Premack and Woodruff (1978)—we have a 'theory' of another mind because mental states cannot be directly observed, but must be inferred from behavior.

The ability to reason about mental states from observed behavior is called mind-reading. Mind-reading and ToM are manifest in humans beyond infancy. They form perhaps the central component of human social intelligence. Whether humans are alone in possessing this ability remains the subject of debate. However, it is commonly accepted that the ToM component is present, and most fully developed, in the human species.

ToM is assumed in game theory without apology. Agents in a game are often automatically assumed to share common knowledge about various aspects of the situation—preferences, the structure of the game, the rationality of the players, etc. The most general cases involve hierarchies of beliefs—beliefs about others' preferences, beliefs about others' beliefs about beliefs, and so on.[45]

Although the suite of cognitive mechanisms supporting ToM appear to be innate in humans, ToM emerges ontogenetically throughout development. An early illustration of this emergence is provided by Wimmer and Perner (1983). They tested children's understanding of two sketches. Each sketch involved a protagonist in a room storing an object in a location, and then leaving the scene, only to have an antagonist move the object to a different location. The child subject observes the surreptitious move of the object, and must know the protagonist could not have observed it. The child then has to indicate where the protagonist would look for the object when returning to the room. Subjects ranged in age from 3 to 9 years. No subject less than 4 years old guessed the correct answer—that the protagonist would look for the object where she hid it. Less than 57% of 4-6 year olds guessed correctly, but 89% of 6-9 years olds guessed correctly.

A second sketch involved the protagonist either helping a friend find the object, or undermining a foe that wants the object. Fully 85% of the older children inferred that the protagonist would lead his friend to where the object was hidden originally, and that he would lead the foe to the other location. Hence, around age 4, there emerges an ability to attribute beliefs to others that differ from one's own, and, moreover, to predict the corresponding behavior.[46]

Baron-Cohen, Leslie and Frith (1983) develop the 'Sally-Ann Test', in which children observe a puppet-play version of the Wimmer and Perner (1983) scenario. In line with Wimmer and Perner, they find that attribution of false belief emerges around age 4. Interestingly, they find that only 20% of subjects diagnosed with

---

[45] Robalino and Robson (2012) discusses the ToM aspects of game theory. Harsanyi (1967, 1968a,b) develops an equilibrium notion that finesses the often unwieldy hierarchical description.

[46] The Wimmer and Perner test—and many others based on it that followed— are known as 'false-belief' tests.

autism give the correct answer in the false-belief task, suggesting that autism involves a ToM deficit.[47].

Following Wimmer and Perner, several attempts have been made to push back the age at which false-belief reasoning emerges. The cutoff has subsequently been the subject of intense debate (Scott and Baillargeon (2017)). For example, Onishi and Baillargeon (2015) test false-belief understanding in infants using a design based on the tendency of infants to gaze longer at events that violate their expectations. Their non-verbal test suggests that 15 month-olds can understand false-beliefs.

Further tests suggest that other components of ToM emerge already in pre-verbal infants. Repacholi and Gopnik (1997), for example, find that, around 18 months, infants begin to differentiate their own desires from the desires of others. In their test, children observe an individual (non-verbally) expressing disgust when tasting one food and pleasure when tasting another food. The children then had an opportunity to offer the individual the two types of food. Whereas 14 month-old subjects offered the food they themselves preferred, 18 month-olds offered the individual the food that induced a happy reaction initially, even when this was their own less preferred option.

There is evidence of cultural variation in several aspects of ToM (Lillard (1998)). But some basic components appear to be universal. For example, Avis and Harris (1991) find that, among the children of the Baka, a group of hunter-gatherers, older children (mean age 5) pass a version of the Sally-Ann test, while most younger children (mean age 3.5) do not.

5.0.1. *ToM in Other Species.* Premack and Woodruff (1978) first asked, “Does the chimpanzee have a theory of mind?” Answers to such questions might help identify the evolutionary preconditions for this ability. Heyes and Frith (2014), for example, argues that the sophisticated expression of ToM that is observed in humans requires uniquely human cognitive byproducts, such as language and culture. ToM in non-verbal species would imply verbal language is not needed for

---

[47] The notion that autism might entail a ToM deficit was later criticized. For example, it has been suggested that the Sally-Ann results with autistic children might result from a language deficit rather than diminished ToM ability (Bloom and German (2000))

this ability, although it might enhance it. For a survey of ToM animal studies see Krupenye and Call (2019). Several species of great apes direct their gaze to where an agent falsely believes an object is hidden, even though the apes know the object is located somewhere else (Krupenye et al. (2016)).

The presence of ToM in non-human animals has been debated since Premack and Woodruff (1978). For non-human species, a test for ToM must be non-verbal. Any display of the ability must also be non-verbal. This makes identification challenging. For example, a deceptive action by an animal— such as hiding food—could be interpreted as evidence of ToM. However, the deceptive strategy might have plausibly arisen via adaptive learning, requiring no reasoning at all about others' mental states. (See Whiten (1996), Heyes (2015), and Krupenye and Call (2019) for discussions of the difficulties involved in ToM animal studies.)

5.0.2. *Hunter Gatherers and ToM.* Whiten and Erdal (2012) argue that ToM is a key pillar in the development of a socio-cognitive niche that facilitated hunter-gatherer social structure, and led to its success and predominance.

Dunbar (1992) argues that human intelligence is driven by the social environment (the social intelligence hypothesis). Group size in primates correlates with relative neocortex volume. Ecological problem solving seems to have played a smaller role. For example, there is no relationship between relative neocortex size and range size. Dunbar (2003) reviews evidence for the social intelligence hypothesis and for alternative explanations of primate intelligence.

5.1. **The Evolution of ToM.** We now address the question of why ToM might have evolved. Although ToM is clearly advantageous in social situations, its supporting cognitive components are likely to be biologically costly (Dunbar (1992), Dunbar (1998)). The question then is , what evolutionary edge does ToM confer over less sophisticated approaches to strategic interaction, such as, for example, adaptive learning?

The basic issue might be illustrated with an example of deception in animals. Scrub-jays prefer to hide food in a distant location when another jay is observing, but are equally likely to hide food at a near or distant location when the observer's view is occluded by an opaque barrier (Dally, Emery and Clayton (2005)). How do we explain this behavior? A ToM based explanation is that the deceptive jay understands her observer's mental states and understands their behavioral implications. The ToM jay hides her food in order to intentionally thwart her observer. Another explanation doesn't involve reasoning about others' mental states at all. The jay might have had her cache raided before, and has learned this happens less often whenever she stores food far from other jays. The hiding strategy is met with success and thus reinforced.

Robalino and Robson (2016) studies the evolution of a key component of ToM, which they christen "theory of preferences" (ToP). ToP refers to the ability to impute preferences to others, to learn these preferences in strategic situations, and to use this knowledge to predict behavior in novel circumstances. The key result is that ToP has a sharp evolutionary edge over naivety when the strategic environment is characterized by persistent strategic novelty. A naive approach, such as reinforcement learning, might do well when novelty is introduced sufficiently slowly. However, agents with ToP can adapt to a higher rate of arrival of novelty than any reasonably specified naive type. The advantage to ToP is that it confers an ability to extrapolate knowledge about preferences across games.

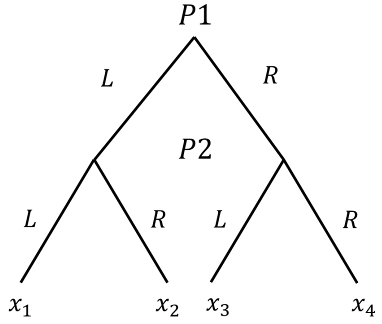

FIGURE 14. The game played by $P1$ and $P2$ in period t. The outcomes x_1, x_2, x_3, x_4 are drawn uniformly from the set of outcomes, X_t.

Consider a dynamic environment where, in each period, a perfect information game is generated by drawing outcomes from a set, X_t, that grows over time. At each date, a large population of individuals is matched to play the realized game. The situation where the game has two players is displayed in Figure 14. Suppose all players know their own preferences over the outcomes, and that they choose a dominant action whenever one is available.[48]

Robalino and Robson (2016) considers a horse-race between two types of individuals in this environment. A naive type is essentially a reinforcement learner. This type learns how to behave in a game through repeated exposure to the specific game. When a game is new, the naive type maps the game to an arbitrary choice. On the other hand, a ToP type knows that her opponents' choices are motivated by preferences, endeavors to learn these preferences, and can extrapolate to novel games what has been learned in previous circumstances.

The naive type is thus characterized by an insensitivity, or blindness, to others' preferences. Such blindness might, nevertheless, appear sophisticated, in a limited sense. In Figure 14, for example, a salient naive strategy for $P1$ is, in each novel game, to choose the action that yields the highest 50-50 average of the payoffs that follows the action. For example, in Figure 14, $P1$ will choose L if

$$u_1(x_1) + u_1(x_2) > u_1(x_3) + u_1(x_4).$$

Actually, this would be the fully Bayesian optimal choice for $P1$ in the first round, given only that $P2$'s preferences are symmetrically distributed, so that the two choices that $P2$ might make after either choice by $P1$ are equally likely. Of course, this rule yields the appropriate choice whenever $P1$ has a dominant action.

The key difference is the rate at which the two types learn after the first round—while the naive type only learns game by game, the ToP type learns preferences. If a ToP type of $P1$ had previously observed $P2$ choose x_1 over x_2, and x_3 over x_4, perhaps in separate games, then she can predict $P2$'s choice in the game from Figure 14 even if this particular game has never arisen before.

---

[48] Suppose there are no super-game effects because in each period the $P1/P2$ pairs are randomly matched.

Robalino and Robson (2016) prove the following. If the set of outcomes grows very slowly, then both the naive type and the ToP type can keep up with the arrival of novel outcomes, and thus both types can learn effectively. If outcomes arrive too fast, then neither type can keep up. Although the ToP is able to learn more efficiently, neither type can keep up with the arrival rate of novelty. In an intermediate range of the arrival rate of novelty, however, the naive type is swamped, essentially facing a new game each period, while the ToP type can acquire her opponent's preferences and thus learn to play appropriately in every circumstance in the long run.

This key intermediate range of the arrival rate of strategic novelty thus provides the necessary impetus for ToP—a low enough arrival rate so the ToP type can learn, but sufficiently high to swamp the naive types. In this range, the ToP type will have a sharp advantage over the naive individuals even if there is a cost to ToP, arising perhaps from the required cognitive resources, provided this cost is sufficiently small. The reason the ToP type can cope with a higher arrival rate is because it needs to track far less information than the naive type. The set of 2-by-2 games, when there are N outcomes is of the order of N^4, while the binary choices the ToP must learn is of the order of N^2.

Kimbrough, Robalino, and Robson (2017) further develop this idea, establishing a precise sense in which the set of 2-by-2 games that can be generated from a set of N outcomes is more complex than the set of complete binary relations over the outcomes, which in turn is more complex than the set of transitive orderings of the outcomes. A sophisticated ToP type that knows her opponents' preferences are transitive will then learn faster than a simple ToP type that learns via binary choices. In turn, simple ToPs will beat a naive type that endeavors to learn games. Kimbrough, Robalino and Robson (2017) further designs and executes an experiment to test the ability to learn others' preferences in one strategic setting and to extrapolate this knowledge to novel games.

As in Section 5.1 consider a two-stage, two-action, perfect information game tree. Player 1 (P1) moves first, then Player 2 (P2). Each player has preferences over a fixed finite set of outcomes, X, and each player knows his own preferences. Let

G be the set of all extensive form games as just described with outcomes from X. Suppose $P1$ and $P2$ are to play a randomly drawn game, $g \in G$. Before $P1$ and $P2$ play the game, however, suppose $P1$ observes a 'training' history h informing $P1$ of $P2$'s subgame perfect strategy in M games, where each of these are independently drawn from G with replacement.

In all of the games $P2$'s problem is straightforward. The question is, how does $P1$ condition his choice in g on the training history h? We consider three types of $P1$, differing in how they incorporate information across games.

***Definition* 5.1.** A naive type of $P1$ makes an arbitrary choice in g if the game is not observed along h. Otherwise, the naive type chooses a best response to the known SPE choices of $P2$ in the game.

As in Robalino and Robson (2016), this naive type considers each game an independent circumstance, and thus requires direct exposure to each particular game in order to learn the appropriate choice. The naive player, by luck, might get it right in a novel game, however, for any reasonable specification of naivete, she will initially make a suboptimal choice in some of the games.

This specification gives the naive types generous credit, in the sense that they need only *one* exposure to a given game to learn how to play it. This only strengthens the results since we will show that, even in this case, there is a marked edge to the ToP types.

A ToP type understands that $P2$'s choices along h are motivated by a complete binary relation over the set of outcomes. If a ToP type observes $P2$ choosing an outcome x when y was available, then she will know that $P2$ prefers x to y, and can predict this choice when it arises in any circumstance. Hence, the ToP type is sure to choose optimally in g whenever both corresponding binary choices of $P2$ have been observed along the training history, even if the game itself has not been seen before.

***Definition* 5.2.** If $P2$'s choice in a subgame of g is revealed along the training history, then the ToP type of $P1$ uses this to formulate her best response in g. Otherwise, the ToP type attributes an arbitrary strategy to $P2$ in the subgame and best responds to this.

The ToP ignores the implications of transitivity, having to directly observe a binary choice in order to infer the corresponding preference. A more sophisticated ToP understands that $P2$ has transitive preferences and can apply this to correctly predict her opponent's choice in a subgame that may not have been seen previously.

***Definition* 5.3.** A ToP* type of $P1$ chooses a best response in g to an arbitrary $P2$ strategy implied by a transitive ordering of X, where this ordering is consistent with the $P2$ strategies observed along the training history.

The key issue is now how effective the various types are in using the training history to learn about G. Clearly, the ToP* extracts more useful information from a given history than does the ToP, and this is also true for the ToP relative to the naive type. For instance, if a ToP learns the appropriate choice in g by observing h, then so must the more sophisticated ToP* that understands transitivity. At the same time there are cases where the ToP* learns but the ToP does not. Intuitively then, the ToP* is the most efficient learner, and the naive type the least.

The edge to sophistication can be made precise by considering a sufficiently large set of games. Recall that h entails M examples of uniformly drawn games from G, and their corresponding $P2$ strategies. Consider a notion of learning in the current circumstance—

***Definition* 5.4.** A type $\tau \in \{\text{naive}, \text{ToP}, \text{ToP}^* \}$ can $(1-\epsilon)$-learn G after M examples if the following holds. For any transitive preference ordering of $P2$, $P1$'s choice in g after observing a randomly drawn history h is the subgame perfect choice with probability no less than $1-\epsilon$, uniformly in the $P2$ preference ordering.

That is, M examples suffice for $P1$ to learn $P2$'s preferences in any G with probability $(1-\epsilon)$. Given an arbitrary preference relation of his opponent, training on M

examples ensures $P1$ will get it right in at least a fraction $(1-\epsilon)$ of the games. This notion of learning —probably approximately correct learning, (PAC-learning)— is due to Valiant (1984)[49]

Clearly, the number of examples required for a $P1$ type to learn about G will depend on the parameter ϵ, however, it will also depend on the complexity of what the type is really trying to learn. PAC-learning theory makes a tight connection between the complexity of the concept being learned and the number of training examples required (see Kalai (2003) for an detailed discussion of the theory). Recall that the naive type is learning a set of games, while the ToP type is learning binary choices, and the ToP* type is learning a linear ordering. The set of games is the most complex, and the transitive ordering the least, in a way that can be made precise in PAC-learning theory. We have the following result, which is proved in Kimbrough, Robalino and Robson (2017)—

Theorem 5.1. *Suppose there are at least* 4 *outcomes. There are positive continuous functions,* $L(\epsilon)$*, and* $U(\epsilon)$ *with* $L(\epsilon) < U(\epsilon)$*, such that the following holds. (1) A naive type of* $P1$ *is sure to* $1-\epsilon$ *learn* G *after* M *examples if* $M \geq |X|^4 \cdot U(\epsilon)$*, but not if* $M < |X|^4 \cdot L(\epsilon)$*. (2) A ToP type of* $P1$ *is sure to* $1-\epsilon$ *learn* G *after* M *examples if* $M \geq |X|^2 \cdot U(\epsilon)$*, but not if* $M < |X|^2 \cdot L(\epsilon)$*. (3) A ToP** *type of* $P1$ *is sure to* $1-\epsilon$ *learn* G *after* M *examples if* $M \geq |X| \cdot U(\epsilon)$*, but not if* $M < |X| \cdot L(\epsilon)$*.*

That is, the ToP* learns the fastest, requiring the least number of examples, and the naive type learns the slowest. The differences in speeds of learning are, moreover, stark, and reveal a sharp edge for ToP* over ToP, and for ToP over the naive type. As in Robalino and Robson (2016), the sophisticated types learn more efficiently because they can extrapolate to a novel game information that was acquired in a different circumstance.

The above results show that an increase in the number of strategic problems might lead to a transition where there is a sharp edge to ToM. This is in line with theory and evidence from other disciplines regarding the evolutionary growth of human intelligence—for example, a central argument there is that the transition from an

---

[49] See Kalai (2003) also, where the the theory is applied to the related problem of the learnability of rational preference orderings.

arboreal niche to the more complex circumstances in the savannah led to an increase in human intelligence. (For discussion see, for example, Robson and Kaplan (2003).)

5.2. **Theory of Preferences: Experimental Evidence.** Kimbrough, Robalino and, Robson (2017) further design an experiment testing the ability to learn others' preferences and to exploit this knowledge strategically. The design is based on the simple theoretical learning model discussed in the previous section.

As in the above theoretical model the experiments involve two-stage two-action games of perfect information. In each experimental session, an even number of subjects are assigned the roles of $P1$ and $P2$, where these roles are fixed throughout the session. There is a fixed set of outcomes, X, where each $x \in X$ induces a unique monetary payoff for each player role, Pi, $i = 1, 2$. All Pi's, $i = 1, 2$, obtain the same monetary payoff from each given $x \in X$. Hence, each subject in a given role has the same induced preferences over X. Each subject knows her own preferences, but not the preferences of players in the other role.

In each period, $t = 1, \dots, T$, of a session, a game is formed by uniformly drawing 4 outcomes from X and appending these to 2-by-2 extensive form. $P1$'s and $P2$'s are then randomly and anonymously matched to play the game. At the end of the period, all the $P1$ subjects observe the choices of all the $P2$s at any nodes that were reached by any $P2$ during the period of play. As in the theoretical model, if a $P1$ subject has ToP, she can learn about $P2$'s preferences. If the likelihood of SPE play by $P1$s varies with the history of play, this might reveal the extent to which subjects have ToP and properly apply it.

The overall trend of play suggests substantial learning. This is displayed in Figure 15 which shows the fraction of $P1$s making the SPE choice in each period. An upward trend indicates learning. The pool of games is restricted over time. The goal was to observe learning but, for some games, the SPE choice could too easily be inferred. Specifically, in each session, after period 15, no game gave the $P1$s a

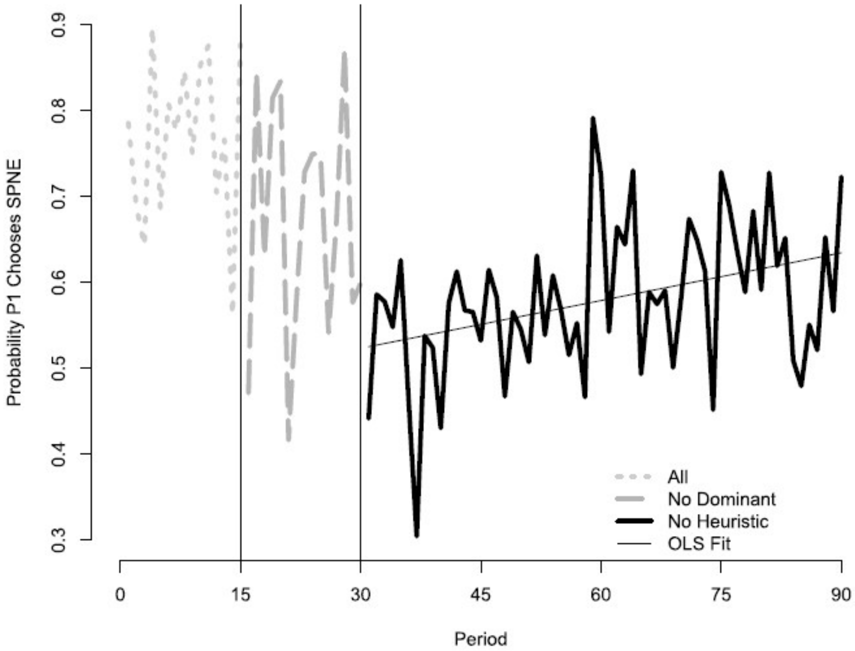

FIGURE 15. Learning an opponent's preferences. This is Figure 2 from Kimbrough, Robalino and Robson (2017). Reuse of this Elsevier material is permitted under STM guidelines.

dominant choice; after period 30, no SPE choice of $P1$ was consistent with a 50-50 heuristic, where this involved choosing the action with the highest 50-50 mean payoff for $P1$ in the resulting subgame.

To determine the type of learning observed—naive, ToP, or ToP*— the probability of SPE choice in a game is regressed on indicators of (1) whether the game was played previously (naive learning), (2) whether the pairwise choices of $P2$ in the game had been observed (ToP learning), (3) whether the pairwise choices of $P2$ in the game are implied by the history of play, given that the $P2$s have transitive preferences (ToP* learning).

Results from two main regressions are presented in Table 1. The first column examines how $P1$'s use the public history of play. The second column examines how

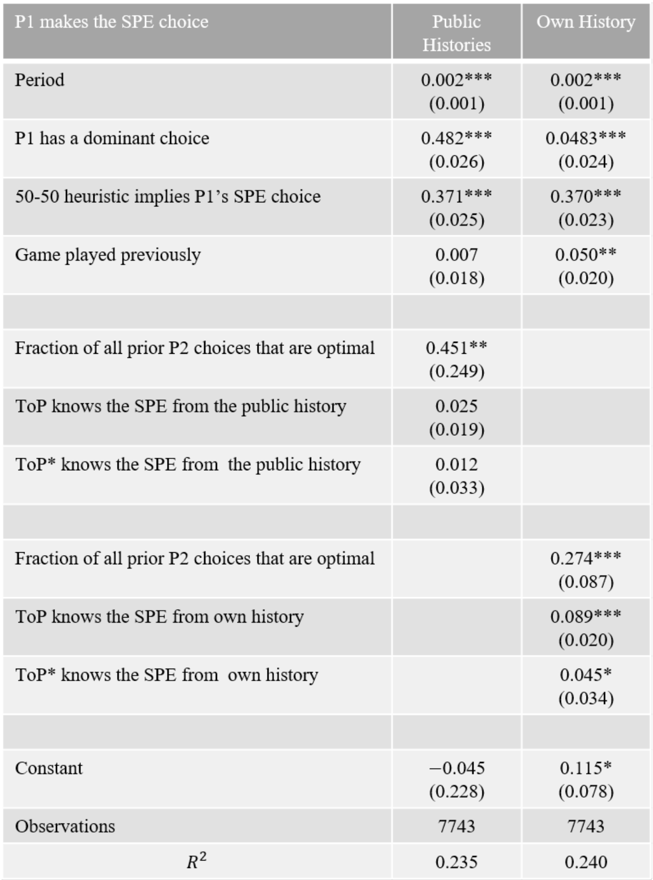

P1 makes the SPE choice	Public Histories	Own History
Period	0.002*** (0.001)	0.002*** (0.001)
P1 has a dominant choice	0.482*** (0.026)	0.0483*** (0.024)
50-50 heuristic implies P1's SPE choice	0.371*** (0.025)	0.370*** (0.023)
Game played previously	0.007 (0.018)	0.050** (0.020)
Fraction of all prior P2 choices that are optimal	0.451** (0.249)	
ToP knows the SPE from the public history	0.025 (0.019)	
ToP* knows the SPE from the public history	0.012 (0.033)	
Fraction of all prior P2 choices that are optimal		0.274*** (0.087)
ToP knows the SPE from own history		0.089*** (0.020)
ToP* knows the SPE from own history		0.045* (0.034)
Constant	−0.045 (0.228)	0.115* (0.078)
Observations	7743	7743
R^2	0.235	0.240

TABLE 1. Learning an opponent's preferences in circumstances that would enable naive, ToP, or ToP* learning. Clustered standard errors in parenthesis. Under a one-tailed hypothesis test, * means $p < 0.10$, ** means $p < 0.05$, and *** means $p < 0.01$. This is an excerpt from Table 1 from Kimbrough, Robalino and Robson (2017). Reuse of this Elsevier material permitted under STM guidelines.

a $P1$ conditions choice on histories involving own opponents only. The results suggest a bias in the use of information. Subjects tend to focus more on their own personal histories, seeming to pay less attention to the public history. ToP learning appears to be better established than ToP* learning. Overall, there is evidence that subjects endeavor to infer others' preferences and that they apply knowledge of preferences appropriately to make optimal choices. Subjects also completed two well-known survey instruments that are used in psychology to measure ToM (Hurley (2007)). Interestingly, ToP ability as measured by the game experiments correlates with responses on the ToM surveys from psychology. Hence, the notion of ToP examined experimentally corresponds with ToM as understood by psychologists.

5.3. **Related Papers on ToM in Economics.**

In addition to the papers discussed above, there is a small literature in economics that covers several aspects of ToM. We now outline some of this relevant work.

5.3.1. *Theoretical Papers.*

A small number of papers examine the evolutionary emergence of strategic intelligence. A key question is whether sophistication is evolutionarily inevitable—that is, does it confer a sufficient advantage over more simple, and perhaps less costly, approaches? Foremost among these is Stahl (1993), which considers a cognitive hierarchy of types, ranging in levels of sophistication. A smart_0 player does not reason about the game at all, automatically choosing an arbitrary strategy instead. For $n \geq 1$, a smart_n type can reason about which choices are n-level rationalizable, and knows that no smart_{n-1} player would use a strategy that is not $(n-1)$-level rationalizable. A key result of Stahl is that, in a stationary strategic environment, it does not always pay to be be smart. In particular, smart_0 players that happen to choose the right strategy cannot be driven out by more sophisticated agents. In some cases, when sophistication entails a cost, the smart_0 will drive the smart types to extinction.

In a similar vein, Mohlin (2012) studies the closely related level-k approach in a more general model allowing for multiple games, learning, and partial observability of type. It remains true, however, that low reasoning types co-exist in the long run with higher types. In a Bayesian game setting, Grauer and Kuzmics (2020) show that, in conflict games, the probability that a player will pay a small cost to acquire her opponent's preferences is bounded away from both zero and one. The result holds given an arbitrary positive cost of sophistication. All of the above works examine a stationary strategic environment. On the other hand, the

Robalino and Robson (2016) result in favor of sophistication is driven by persistent strategic novelty.

5.3.2. *Experimental Papers.* On the experimental front there is a large literature examining learning in games. Although deviations from classic rationality are well-known and often observed, there is substantial evidence that learning is more sophisticated than in simple adaptive learning. An important paper is Erev and Roth (1998) which compares the predictive performance of various adaptive learning models. They show that a simple reinforcement model fits the data better than assuming equilibrium is instantly reached. An even better fit is attained when experimentation and forgetting are added to the basic setup. A model where players condition on beliefs about others' strategies yet further improves predictive performance, suggesting individuals reason about other players, as in ToM.

Camerer et al. (2002) develop an elaboration of reinforcement learning called experience weighted attraction (EWA). EWA combines elements of reinforcement learning with the more sophisticated fictitious play. The basic model is extended to incorporate sophisticated learning, where a player understands her opponent might be adaptive, or perhaps, may likewise be a sophisticated learner. Such sophistication improves predictive power relative to an EWA model with only unsophisticated agents.

Knoepfle, Camerer, and Wang (2009) use eye-tracking technology to show that subjects examine others' payoffs carefully, further supporting elements of strategic sophistication. In another study using eye-tracking, Polonio, Di Guida and Coricelli (2015) show that there is heterogeneity in the pattern of information acquisition used by subjects in games, but that subjects' patterns are stable across games, and that these patterns predict behavior. Some of these information acquisition patterns correspond to ones discussed here. For example, some individuals in the study systematically neglect the payoffs of their opponents like our naive individuals. Interestingly, these individuals' eye-movements are consistent with searching for a dominant choice, or the strategy yielding the highest mean payoff. Other subjects performed saccades that alternated fixation on own and others' payoffs, appearing to consider the rationality of their opponents, akin to our ToP types. Related is the question of whether subjects are even aware they are in a strategic situation, and that their choices ought to depend on what they believe about others. Fehr and Huck (2016) show that some subjects are strategically unaware, and that awareness is correlated with cognitive ability.

A key element of ToM is recursive rationality. Although, in principle, there is no limit on the number of recursive embeddings that are possible, in practice, people most often operate at one or two levels (Camerer, Ho, and Chong (2004), Bosch-Domenech et al (2002)). Kneeland (2015) examines the prevalence and depth of recursive rationality in experimental subjects using a novel application of ring games. In a ring game, player 1's payoff depends on player 2's choice, player 2's payoff depends on player 3's choice, $, \ldots,$ $(n-1)$'s payoff depends on player n's choice; and player n's payoff depends on player 1's choice. This ring structure allows higher order payoff to vary independently of lower order payoffs, and thus allows the identification of an individual's depth of recursive rationality. About 6 percent of the subjects are not rational at all. The remaining pool of subjects consists of first-order through fourth-order rationals, with approximately equal shares of each.

In an interesting study involving children, Fe, Gill, and Prowse (2022) find that level-*k* depth correlates with both ToM ability and cognitive ability. Older children incorporate information about their opponents' cognitive ability in a manner suggestive of strategic sophistication. There is some evidence, however, that observed depth of sophistication, as in level-*k*, depends on the circumstance. Georganas et al. (2015), for example, finds in a test of level-*k* that an individual's level of reasoning will vary across games. Moreover, it was not possible to rank subjects by depth of reasoning—the most sophisticated subjects in one circumstance were not always the most sophisticated in another case.

6. FUTURE DIRECTIONS

Of all the topics covered in this chapter, perhaps the leading candidate for extension is theory of mind. All that was sketched here was one aspect of that—the imputation of utility to others as relevant in particular sequential strategic interactions. This merely gathers the lowest-hanging fruit concerned with this issue.

Firstly, it would be worth examining how theory of mind could be evolutionarily advantageous in other strategic interactions. A leading example here would be simultaneous-move games. These would again involve imputing utility to others but also considering the evolutionary advantages of understanding that an opponent would not choose a dominated strategy, as the first step in iterated elimination of dominated strategies. Such games could be examined in the laboratory.

Another interesting aspect of theory of mind is the imputation of beliefs to others as also relevant in strategic interactions, especially when beliefs have effects distinct from those of utility.

Promising empirical work derived from this chapter concerns the transmission of a genetic component to time preference. How would the ancestral environment of various countries shape the time preference seen there now? Section 3.2.6 argues that low ancestral mortality and low ancestral fertility might be reflected in high modern patience. If ancestral mortality and fertility for the last few thousand years could be estimated, the prediction could be checked against the data in Falk, Becker, Dohmen, Enke, Huffman, and Sunde (2018, Figure 1, for example).

From a big picture viewpoint, it seems desirable for future research to develop and test integrated genetic and cultural models for the transmission of economic preferences. Such dual transmission mechanisms could feature the possibility that copying other individuals promoted evolutionary success. The coevolution of genes and culture is intrinsically plausible. The reason for confining attention here at first to purely genetic inheritance is for parsimony and simplicity but also as part of a conservative methodological approach. Biological evolution is the obvious candidate for justifying the conventional assumption in economics that preferences are exogenous and fixed, at least in the short run.

7. CONCLUSIONS

Staying within the ambit of genetically based preferences, the literature surveyed in this chapter still reflects a deliberate reluctance to make wholesale changes to economic theory. For time preference as considered in Sections 3.2-3.2.6, for example, the central model of intertemporal preferences is a standard model based on additively separable flow utility, mortality and a constant discount factor. The novel twist is just that this criterion is a result rather than an assumption, where a constant component of the discount factor deriving from population growth is augmented by mortality. It is argued that such a discount factor might have originated in hunter-gatherer societies despite the apparent paucity of opportunities to make large intertemporal tradeoffs and that this discount factor might be appreciable despite the arithmetic impossibility of substantial growth in human populations over the long run.

Section 4 describes how the predictions of an evolutionary approach based on demography can come from left field, despite the reluctance to make wholesale changes. These predictions are that individuals should be more averse to aggregate risk than to idiosyncratic risk. If such left field predictions are borne out anyway, they boost greatly the inherent appeal of the approach. We describe work by Oprea and Robalino (2024) that corroborate that such a distinction can be seen in experiments.

Section 2 adopted a different evolutionary approach. Rather than considering a Darwinian battle for the numerical supremacy of various rates of time preference and various attitudes to risk, we draw inspiration from direct neuroscience observations. We sketch how utility is apparently directly represented as noisy neural firing rates, so creating the need for utility to adapt to where the heavy traffic is. We present a model of adaptive utility and show it makes predictions—about a preference for growth, for example—that are substantiated. We delve into the possibility that utility might further be reflected as pleasure which then provides an explanation for the hedonic treadmill and raises thorny questions about happiness economics.

These results demonstrate the power of an evolutionary approach to buttress economic theory, to suggest novel hypotheses, to extend its application. The link to neuroscience, in particular, holds out the promise of eventually better integrating economics and other sciences that study the same human beings.

REFERENCES

[1] Jeremy Avis and Paul L. Harris. Belief-desire reasoning among Baka children: Evidence for a universal conception of mind. *Child Development*, 62(3):460–467, 1991.

[2] Simon Baron-Cohen, Alan M. Leslie, and Uta Frith. Does the autistic child have a "theory of mind"? *Cognition*, 21(1):37–46, 1983.

[3] Anke Becker, Benjamin Enke, and Armin Falk. Ancient origins of the global variation in economic preferences. In *AEA Papers and Proceedings*, volume 110, pages 319–323, 2020.

[4] Gary S. Becker. *Human Capital*. Columbia University Press, New York, 1964.

[5] Jeremy Bentham. *An Introduction to the Principles of Morals and Legislation*. Dover Publications, Mineola, New York, 2007.

[6] Kent C Berridge. The debate over dopamine's role in reward: the case for incentive salience. *Psychopharmacology*, 191:391–431, 2007.

[7] Kent C. Berridge. Wanting and liking: Observations from the neuroscience and psychology laboratory. *Inquiry*, 52(4):378–398, 2009.

[8] Kent C. Berridge and Morten L. Kringelbach. Pleasure systems in the brain. *Neuron*, 86(3):646–664, 2015.

[9] Kenneth George Binmore. *Game theory and the social contract volume I: Playing fair*, volume 1. MIT Press, 1994.

[10] Paul Bloom and Tim P. German. Two reasons to abandon the false belief task as a test of theory of mind. *Cognition*, 77:B25±B31, 2000.

[11] Antoni Bosch-Domènech, Josè G. Montalvo, Rosemarie Nagel, and Albert Satorra. One, two, (three), infinity, ...: Newspaper and lab beauty-contest experiments. *American Economic Review*, 92(5):1687–1701, 2002.

[12] James J. Bull. Sex determination in reptiles. *The Quarterly Review of Biology*, 55(1):3–21, 1980.

[13] Colin F. Camerer, Teck-Hua Ho, and Juin-Kuan Chong. Sophisticated experience-weighted attraction learning and strategic teaching in repeated games. *Journal of Economic Theory*, 104(1):137–188, 2002.

[14] Colin F. Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119(3):861–898, 2004.

[15] John Y Campbell and John H Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251, 1999.

[16] Andrew E. Clark, Sarah Flèche, Richard Layard, Nattavudh Powdthavee, and George Ward. *The Origins of Happiness: The Science of Well-Being over the Life Course*. Princeton University Press, Princeton, 2018.

[17] Kauê Machado Costa and Geoffrey Schoenbaum. Dopamine. *Current Biology*, 32(15):R817–R824, 2022.

[18] Paolo Crosetto and Antonio Filippin. The "bomb" risk elicitation task. *Journal of Risk and Uncertainty*, 47(1):31–65, 2013.

[19] Joanna M. Dally, Nathan J. Emery, and Nicola S. Clayton. Cache protection strategies by western scrub-jays, aphelocoma californica: Implications for social cognition. *Animal Behavior*, 70(6):1251–1263, 2005.

[20] Susanne Dreborg, Görel Sundström, Tomas A. Larsson, and Dan Larhammar. Evolution of vertebrate opioid receptors. *Proceedings of the National Academy of Sciences*, 105(40):15487–15492, 2008.

[21] James S. Duesenberry. *Income, Saving, and the Theory of Consumer Behavior*. Cambridge University Press, Cambridge, 1949.

[22] R. I. M. Dunbar. Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*, 22(6):469–493, 1992.

[23] R. I. M. Dunbar. The social brain hypothesis. *Evolutionary Anthropology*, 6(5):178–190, 1998.

[24] R. I. M. Dunbar. The social brain: Mind, language, and society in evolutionary perspective. *Annual Review of Anthropology*, 32(1):163–181, 2003.

[25] Christian Dustmann and Costas Meghir. Wages, experience and seniority. *Review of Economic Studies*, 72(1):77–108, 2005.

[26] Richard A. Easterlin. Does economic growth improve the human lot? Some empirical evidence. *Nations and Households in Economic Growth: Essays in Honor of Moses Abramovitz*, 1974.

[27] Richard A. Easterlin and Kelsey J. O'Connor. The Easterlin paradox. In *Handbook of labor, human resources and population economics*, pages 1–25. Springer, 2022.

[28] Larry Epstein and Stephen Tanny. Increased generalized correlation: A definition and some economic consequences. *Canadian Journal of Economics*, 13(1):16–34, 1980.

[29] Ido Erev and Alvin Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 88(4):848–881, 1998.

[30] Armin Falk, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde. Global evidence on economic preferences. *Quarterly Journal of Economics*, 133(4):1645–1692, 2018.

[31] Eduardo Fe, David Gill, and Victoria Prowse. Cognitive skills, strategic sophistication, and life outcomes. *Journal of Political Economy*, 130(10):2643–2704, 2022.

[32] Dietmar Fehr and Steffen Huck. Who knows it is a game? on strategic awareness and cognitive ability. *Experimental Economics*, 19(4):713–726, 2016.

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

[34] Robert H. Frank and Robert M. Hutchens. Wages, seniority, and the demand for rising consumption profiles. *Journal of Economic Behavior & Organization*, 21(3):251–276, 1993.

[35] James C. Frauenthal. Analysis of age-structure models. In *Mathematical Ecology: An Introduction*, pages 117–147. Springer, 1986.

[36] William J Freed. *Motivation and Desire: A New Way to Think about why We Do Everything and Its Basis in Neuroscience*. Springer Nature, 2022.

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

[38] Oded Galor and Ömer Özak. The agricultural origins of time preference. *American Economic Review*, 106(10):3064–3103, 2016.

[39] Florian Gauer and Christoph Kuzmics. Cognitive empathy in conflict situations. *International Economic Review*, 61(4):1659–1678, 2020.

[40] Sotiris Georganas, Paul J. Healy, and Roberto A. Weber. On the persistence of strategic sophistication. *Journal of Economic Theory*, 159:369–400, 2015.

[41] Paul W Glimcher. *Foundations of neuroeconomic analysis*. Oxford University Press, New York, 2011.

[42] Glenn W. Harrison, Morten I. Lau, and E. Elisabet Rutström. Estimating risk attitudes in Denmark: A field experiment. *Scandinavian Journal of Economics*, 109(2):341–368, 2007.

[43] John C. Harsanyi. Games with incomplete information played by 'bayesian' players, i-iii. *Management Science*, 14:159–182, 320–334, 486–502, 1967-68.

[44] Yuval Heller and Ilan Nehama. Evolutionary foundation for heterogeneity in risk aversion. *Journal of Economic Theory*, 2008:1–16, 2023.

[45] Cecilia M. Heyes. Animal mindreading: What's the problem? *Psychonomic Bulletin and Review*, 22:313–327, 2015.

[46] Cecilia M. Heyes and Chris D. Frith. The cultural evolution of mind reading. *Science*, 344(6190):1357, 2014.

[47] Kim Hill and A. Magdalena Hurtado. *Ache Life History: The Ecology and Demography of a Foraging People*. Aldine de Gruyter, New York, 1996.

[48] Wangjie Hu, Ziqian Hao, Pengyuan Du, Fabio Di Vincenzo, Giorgio Manzi, Jialong Cui, Yun-Xin Fu, Yi-Hsuan Pan, and Haipeng Li. Genomic inference of a severe human bottleneck during the early to middle pleistocene transition. *Science*, 381(6661):979–984, 2023.

[49] Robert S.E. Hurley, Molly Losh, Morgan Parlier, Stephen J. Reznick, and Joseph Piven. The broad autism phenotype questionnaire. *Journal of Autism and Developmental Disorders*, 37(9):1679–1690, 2007.

[50] Daniel Kahneman, Peter P Wakker, and Rakesh Sarin. Back to Bentham? Explorations of experienced utility. *Quarterly Journal of Economics*, 112(2):375–406, 1997.

[51] Gil Kalai. Learnability and rationality of choice. *Journal of Economic Theory*, 113(1):104–117, 2003.

[52] Hillard S. Kaplan and Arthur J. Robson. We age because we grow. *Proceedings of the Royal Society B: Biological Sciences*, 276(1663):1837–1844, 2009.

[53] Erik O. Kimbrough, Nikolaus Robalino, and Arthur J.Robson. Applying "theory of mind": Theory and experiments. *Games and Economic Behavior*, 106(1):209–226, 2017.

[54] Terri Kneeland. Identifying higher-order rationality. *Econometrica*, 83(5):2065–2079, 2015.

[55] Daniel T. Knoepfle, Colin F. Camerer, and Joseph T. Wang. Studying learning in games using eye-tracking. *Journal of the European Economic Association*, 7(2-3):388–398, 2009.

[56] Philip A. Kragel, Michael T. Treadway, Roee Admon, Diego A. Pizzagalli, and Emma C. Hahn. A mesocorticolimbic signature of pleasure in the human brain. *Nature Human Behaviour*, pages 1–12, 2023.

[57] Christopher Krupenye and Josep Call. Theory of mind in animals: Current and future directions. *Wiley Interdisciplinary Reviews in Cognitive Science*, 10(6):1–25, 2019.

[58] Christopher Krupenye, Fumihiro Kano, Satoshi Hirata, Josep Call, and Michael Tomasello. Great apes anticipate that other individuals will act according to false beliefs. *Science*, 354(6308):110–114, 2016.

[59] Angeline Lillard. Ethnopsychologies: Cultural variations in theories of mind. *Psychological Bulletin*, 123(1):3–32, 1998.

[60] Andreu Mas Colell, Michael Whinston, and Jerry Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

[61] Erik Mohlin. Evolution of theories of mind. *Games and Economic Behavior*, 75(1):299–318, 2012.

[62] Nick Netzer, Arthur Robson, Jakub Steiner, and Pavel Kocourek. Risk perception: Measurement and aggregation. *Journal of the European Economic Association*, 2024. forthcoming.

[63] William D. Nordhaus. *Managing the global commons: the economics of climate change*, volume 31. MIT press Cambridge, MA, 1994.

[64] Kristine H. Onishi and René Baillargeon. Do 15-month-old infants understand false beliefs? *Science*, 308(8):255–258, 2005.

[65] Ryan Oprea. Decisions under risk are decisions under complexity. *American Economic Review*, 2024. forthcoming.

[66] Ryan Oprea and Nikolaus Robalino. Are humans more averse to aggregate risk? testing an evolutionary economic theory. Working Paper, 2024.

[67] Luca Polonio, Sibilla Di Guidac, and Giorgio Coricelli. Strategic sophistication and attention in games: An eye-tracking study. *Games and Economic Behavior*, 94(1):80–96, 2015.

[68] David Premack and Guy Woodruff. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1(4):515–526, 1978.

[69] Debraj Ray and Arthur J. Robson. Status, intertemporal choice, and risk taking. *Econometrica*, 80(4):1505–1531, 2012.

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

[71] Denis Réale, Andrew G McAdam, Stan Boutin, and Dominique Berteaux. Genetic and plastic responses of a northern mammal to climate change. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(1515):591–596, 2003.

[72] Betty M. Repacholi and Alison Gopnik. Early reasoning about desires: Evidence from 14- and 18-month-olds. *Developmental Psychology*, 33(1):12–21, 1997.

[73] Nikolaus Robalino and Arthur J. Robson. The evolution of strategic sophistication. *American Economic Review*, 106(4):1046–1072, 2016.

[74] Roberto Robatto and Balázs Szentes. On the biological foundation of risk preferences. *Journal of Economic Theory*, 172:410–422, 2017.

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

[76] Arthur J. Robson. The evolution of attitudes to risk: Lottery tickets and relative wealth. *Games and Economic Behavior*, 14(2):190–207, 1996.

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

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

[79] Arthur J. Robson and Hillard S. Kaplan. The evolution of human life expectancy and intelligence in hunter-gatherer economies. *American Economic Review*, 93(1):150–169, 2003.

[80] Arthur J. Robson and Hillard S. Kaplan. The economics of hunter-gatherer societies and the evolution of human characteristics. *Canadian Journal of Economics/Revue canadienne d'économique*, 39(2):375–398, 2006.

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

[82] Arthur J. Robson and Larry Samuelson. The evolution of time preference with aggregate uncertainty. *American Economic Review*, 99(5):1925–1953, 2009.

[83] Arthur J. Robson and Larry Samuelson. The evolution of decision and experienced utilities. *Theoretical Economics*, 6(3):311–339, 2011.

[84] Arthur J. Robson and Larry Samuelson. Evolved attitudes to idiosyncratic and aggregate risk in age-structured populations. *Journal of Economic Theory*, 181:44–81, 2019.

[85] Arthur J. Robson and Larry Samuelson. The evolution of risk attitudes with fertility thresholds. *Journal of Economic Theory*, 205:105552, 2022.

[86] Arthur J. Robson and Larry Samuelson. The evolutionary foundations of preferences. Econometric Society Monograph Series forthcoming, 2026.

[87] Arthur J. Robson, Lorne A. Whitehead, and Nikolaus Robalino. Adaptive utility. *Journal of Economic Behavior & Organization*, 211:60–81, 2023.

[88] Leonard J. Savage. *The Foundations of Statistics*. John Wiley and Sons, London, 1954.

[89] David A. Schkade and Daniel Kahneman. Does living in California make people happy? A focusing illusion in judgments of life satisfaction. *Psychological Science*, 9(5):340–346, 1998.

[90] Wolfram Schultz. Neuronal reward and decision signals: from theories to data. *Physiological reviews*, 95(3):853–951, 2015.

[91] Rose M. Scott and Renee Baillargeon. Early false-belief understanding. *Trends in Cognitive Science*, 21(4):237–249, 2017.

[92] Dale O. Stahl. Evolution of smart$_n$ players. *Games and Economic Behavior*, 5(4):604–617, 1993.

[93] William R. Stauffer, Armin Lak, and Wolfram Schultz. Dopamine reward prediction error responses reflect marginal utility. *Current Biology*, 24(21):2491–2450, 2014.

[94] Nicholas Stern. *Stern Review On The Economics of Climate Change*. Her Majesty's Treasury, London, 2006.

[95] Jetro J. Tuulari, Lauri Tuominen, Femke E. de Boer, Jussi Hirvonen, Semi Helin, Pirjo Nuutila, and Lauri Nummenmaa. Feeding releases endogenous opioids in humans. *Journal of Neuroscience*, 37(34):8284–8291, 2017.

[96] L. G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.

[97] Thorstein Veblen. *The Theory of the Leisure Class*. Routlege, New York, 2017.

[98] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944.

[99] Andrew Whiten. When does smart behavior-reading become mind-reading. In P. Carruthers and P.K. Smith, editors, *Theories of Theories of Mind*, pages 277–292. Cambridge University Press, Cambridge, UK, 1996.

[100] Andrew Whiten and David Erdal. The human socio-cognitive niche and its evolutionary origins. *Philosophical Transactions of the Royal Society B. Biological Sciences*, 367(1599):2119–2129, 2012.

[101] Heinz Wimmer and Josef Perner. Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, 13(1):103–128, 1983.