

Evolutionary Behavioral Economics

Terence C. Burnham, Stephen E. G. Lea,
Adrian V. Bell, Herbert Gintis, Paul W. Glimcher,
Robert Kurzban, Leonhard Lades, Kevin McCabe,
Karthik Panchanathan, Miriam Teschl, and Ulrich Witt

Abstract

This chapter explores how the economic model of individual behavior can be improved by incorporating a number of insights from evolutionary theory and complex systems theory. Insights from psychology, the neurosciences, and the behavioral strand of economics may be better understood from an evolutionary and complexity perspective. It takes an integrated interdisciplinary approach to economic phenomena. Core concepts in economic theory (preference and choice) are clarified and Tinbergen's "four questions" about the origins of behavior are used to provide a framework. Informed by Tinbergen, areas from behavioral science are presented which may be useful for understanding economic behavior: some are directly evolutionary, while others come from scientific contexts informed by evolutionary theory. Each area has yielded well-researched ideas that provide considerable insight into human nature. It concludes with a review of where research stands today and where it could be directed in the future.

Introduction

What kind of modification of the economic theory of individual decision making do we need, and why? We began this project with a belief that the economic model of individual behavior can be improved by incorporating a number of insights from evolutionary theory and complex systems theory. This includes, but is not restricted to, grounding individual behavior in more psychologically and socially informed decision making, along lines already followed in experimental and behavioral economics. Our view, however, is that insights newly gained from psychology, the neurosciences, and the behavioral strand

of economics may be better understood from an evolutionary and complexity perspective.

Why do we think this is important? The fundamental reason is that it is one thing to understand what determines certain small-scale choices and to make, if possible, accurate predictions for specific policy purposes, for example, but it is quite another to map out the long-term causes and consequences (evolution) and connections (complexity) of human behavior in the context of their social and natural environment. Only such a comprehensive understanding will enable us to say whether human beings, and our collective behavior, are on a constructive or a disastrous trajectory.

In our group, we generally agreed that human behavior has, in some sense, lost contact with the natural and ecological constraints that shape it. Some members of the group, however, went further and took the view that this creates a dangerous situation for the sustainability of the planet. If, as these members argue, the “house is on fire,” then insights from evolutionary and complexity theory should help us understand, for example, which motivations propel us toward consumption behavior that adds fuel to those flames. Now more than ever, economists accustomed to modeling social phenomena, and with them other scientists from an array of disciplines (e.g., anthropology, archeology, biology, history, sociology, philosophy, and psychology), are challenged to take up the unique opportunity to contribute collectively to a broader understanding of those phenomena in an effort to make human life on this planet sustainable. It goes without saying that such interdisciplinary cooperation should also be applied to other pressing social phenomena, such as the rise and spread of fundamentalism and, possibly connected with this, the consequences and remedies of social imbalances due to growing inequalities in wealth and access to resources, as well as the need for cultural integration and education.

This chapter is structured as follows. As with all interdisciplinary research, its success depends on a common understanding of the main concepts used. Since the aim of this Forum was to take an integrated interdisciplinary approach to economic phenomena, the concepts that most need clarification are those that lie at the heart of economic theory: preference and choice. Thus, we begin with a brief discussion of these concepts before turning to Tinbergen’s “four questions” about the origins of behavior (Tinbergen 1963). Informed by Tinbergen, we structure a list of ideas from behavioral science that may be useful for understanding economic behavior: some of these are directly evolutionary, while others come from scientific contexts that are informed by evolutionary theory. Each of the areas described has yielded well-researched ideas that can give considerable insight into human nature (though it certainly does not capture them all). From there we review where we are today and where we might want to direct research in the future. The creation of an evolutionary behavioral economics will take many years as well as the concerted effort of many scholars. Thus our discussion necessarily takes an aspirational view.

Preference, Choice, and Decision Making

A widely used textbook (Mas-Colell et al. 1995) describes the mainstream or standard economic model as it was developed originally by Léon Walras, Vilfredo Pareto, John Hicks, Paul Samuelson, Milton Friedman, and others, including Leonard Savage. This standard model generally consists of two possible approaches to individual decision making: one based on preferences, the other on choice. The creators of this model did not intend for the model to be a perfect reflection of human nature. For example, the transitivity assumption assumes there are no cycles in preferences; however, in some circumstances, people do express preferences with cycles. Their motive was practical in nature: the standard model created a mathematical simplification which (as developed later by Abraham Wald and Gérard Debreu) yields tremendous analytic power.

The preference-based approach starts with the assumption that a decision maker has given, stable preferences that are complete and transitive and, on this basis, studies the behavior which follows. Completeness means that the decision maker is able to compare any two alternatives. “It takes work and serious reflection to find out one’s own preferences. The completeness axiom says that this task has taken place: our decision makers make only meditated choices” (Mas-Colell et al. 1995:6). Transitivity rules out the possibility that a sequence of pairwise choices would lead to a cycle. Obviously, Mas-Colell et al. realize that cycles are possible, because they admit that “[s]ubstantial portions of economic theory would not survive if economic agents could not be assumed to have transitive preferences” (Mas-Colell et al. 1995:7).¹

As it happens, if preferences are complete and transitive (i.e., rational according to Mas-Colell et al. 1995) and satisfy continuity, then they can be represented in terms of a continuous utility function. Given constraints, the decision maker is then assumed to optimize or maximize this utility function, which means that she is assumed to choose the best available option out of her affordable opportunity set. Thus, the utility function and the preferences they capture lie at the heart of the preference-based approach to decision making.

In contrast with this preference approach, the choice-based approach takes choices as the primitive feature and imposes a structure on them in terms of a consistency axiom, such as the *weak axiom of revealed preferences* (WARP). This implies that if alternative x has been chosen when alternative y was also available, then there is no situation (or no other budget set) in which y is chosen

¹ Mas-Colell et al. do not define preferences explicitly, but they seem to take them as “feelings of goodness.” This becomes clear in the following sentence, which describes intransitivity or a cycle in pairwise choices: “for example, feeling that an apple is at least as good as a banana and that a banana is at least as good as an orange, but then also preferring an orange over an apple” (Mas-Colell et al. 1995:7).

and x is not. Stated otherwise, in such a case, x is revealed preferred to y .² The choice-based approach “leaves room, in principle, for more general forms of behavior than is possible with the preference-based approach” (Mas-Colell et al. 1995:5). This means that assumptions do not need to be made about the reasons or motives for which the decision maker would choose x over y ; no introspection is necessary; all that is required is that the decision maker acts consistently. This also implies that revealed preferences do not need to be preferences at all, as the reasons for choice may be, in principle, based on something other than the decision maker’s own tastes or feelings of goodness. As it turns out, however, there are certain conditions under which revealed preferences are preferences and vice versa (Houthakker 1950; Uzawa 1960); that is, complete and transitive preferences generate a choice structure that satisfies the weak axiom, and a choice structure that satisfies the weak axiom (provided that enough subsets and choices are available) leads to a rational preference ordering that rationalizes the choices relative to the set of budget sets. Under these conditions, therefore, “[p]references explain behavior; we can interpret the decision maker’s choices *as if* she were a preference maximizer” (Mas-Colell et al. 1995; italics added).

So, choices can be explained by preferences after all. It is for this reason that most economists take choices to be largely equivalent to preferences, even though underlying choices need not be derived from preferences at all, but may occur for other reasons. Although the choice-based approach is preferred by most economists, because “it makes clear that the theory of individual decision making need not be based on a process of introspection but can be given an entirely behavioral foundation” (Mas-Colell et al. 1995:5), this equivalence explains the persistence of the preference-based approach. Mas-Colell, Whinston, and Green themselves write that “[t]his preference-based approach is the more traditional of the two, and it is the one that we emphasize” (Mas-Colell et al. 1995:5). It is also a basis of *as-if* reasoning so widespread in economics. However, two words of caution are required here. First, Mas-Colell et al.’s statement that the decision maker’s choices can be seen *as if* she were a preference maximizer is, strictly speaking, too strong. Here, *as if* refers to the fact that consistent choices can be seen as being rationalizable by an underlying rational preference ordering. Choosing the best available alternative is an

² WARP is a sufficient but not necessary condition to find a preference ordering that rationalizes the choice data as backward. As far as textbook economics is concerned, a sufficient and necessary condition is SARP (S for strong) or more generally GARP (G for generalized). However, those axioms rely on “indirectly revealed preferred” choices, that is, on choices which may, strictly speaking, not have been observed (the core idea of revealed preference theory) but they help to generate empirically useful data. In fact, Mas-Colell et al. (1995:91) mention that the choice-based theory of demand is essentially equivalent to the preference-based theory as SARP imposes transitivity on observed behavior from the outset. This is of course a particular weakness of the choice-based approach: it is one thing to assume transitivity of preferences, but one cannot strictly speaking, impose transitivity on observed choices. See also Wong (2006).

additional assumption about the choice rule of the decision maker, which is not always spelled out clearly in economic textbooks. It goes without saying that there could, in principle, be other choice rules, but maximization is the one favored by most economists.³ Second, a more general interpretation of *as-if* reasoning exists in economics, based on Milton Friedman's version, which has nothing to do with consistent choices but instead refers to the assumption of optimization. Friedman famously stated that the choice rules in decision processes do not need to be descriptively realistic. Even though business people would probably state that they price at average cost with deviations when markets make it necessary, their decisions can legitimately be represented "*as if* they were seeking rationally to maximize their expected returns (generally if misleadingly called "profits") and had full knowledge of the data needed to succeed in this attempt" (Friedman 1953:21).

Economic models that begin with a given utility function and the assumption of optimization do not clarify whether the utility function is appropriate because the underlying behavior is consistent (the *as-if* assumption in the sense of assuming that choices are made as if they derived from preferences), or whether utility maximization is taken to be a model detached from realistic assumptions but still capable of making useful predictions, including at the aggregate level (the *as-if* assumption in Friedman's sense). This is certainly confusing and can be unhelpful in an interdisciplinary context. Indeed, many within our group were uneasy with the *as-if* approach, especially Friedman's version of it. The relationship between models, their assumptions, and the conclusions drawn from them is an important methodological issue that needs to be clarified. In particular, the conclusion that Friedman reached—that models based on heavily simplified assumptions about individual choice will make adequate predictions in the aggregate—needs to be reviewed, as we have a great deal more information about individual choice, and much better ways of studying it, than was available in 1953 (cf. Wilson 2012).

Equally, the assumption of given and stable preferences needs to be reconsidered, as well as the idea of consistent choices. In any case, it needs to be clarified, in much more detail, under which conditions typical economic textbook assumptions do hold, and when it would be better to turn toward a new theory(ies) of behavior and choice, and how behavior changes with motivations, depending on the context and institutions that frame those choices. Especially experimental, behavioral, and neuroeconomics but also other fields such as anthropology and animal studies can help, together with appropriate methodological and epistemological reflections, make economics become an interdisciplinary field of study firmly engrained within the social sciences, such as it should be in the eyes of many members or our group.

³ Hausman (2011) calls this additional assumption the axiom of "choice determination."

Evolutionary Concepts That Might Improve Economics

To further the debate about the model of individual choice needed to construct realistic and valid models of the economy, our discussion focused on specific factors that influence, determine, and contribute to economic decision making, and the means by which they are currently investigated. In a sense, this approach creates a toolbox for the research required to construct a complex, evolutionary account of economic phenomena. Although we have attempted to include as many aspects as possible, the toolbox reflects our own research foci and experience and is incomplete. Thus, we encourage others to look beyond conventional economic research questions, models, and methodologies and to take the tools we suggest on board for a better understanding of human behavior and its influence on its ecology and vice versa.

Tinbergen's Four Questions

The ethologist Niko Tinbergen (1963) argued that a full understanding of any behavior requires us to be able to answer four questions. For convenience, we have reordered them here:

1. **Causation:** what is the mechanism underlying the behavior? This refers both to the immediate stimulus and motivational situation eliciting the behavior, and to the physiological processes (neural, hormonal, muscular) that enable it to occur.
2. **Ontogeny:** what past events in the life of the organism have contributed to the behavior occurring? Of particular concern here are developmental considerations; a person's economic behavior can be profoundly influenced by experiences they had during childhood, or before some key transition such as migration.
3. **Evolution,** by which Tinbergen means evolutionary history: to what extent is the behavior a function of what sort of animal the organism is? Phylogenetically, among animals humans are vertebrates, among vertebrates we are mammals, and among mammals we are great apes. At least some of our economic behavior may be best understood in terms of our vertebrate, mammalian or ape heritage.
4. **Survival value:** how does the behavior contribute to the organism's capacity to survive as an individual and pass on its genes to successive generations? In the modern context, with an understanding of gene-culture evolution and multilevel selection, we need also to consider how the behavior may contribute to the passing on of ideas within a culture and the survival of the economic group.

These questions have been used to help organize knowledge about many different kinds of behavior, including economic behavior (Wilson and Gowdy 2013). The first two questions ask for a proximate explanation of behavior

(how does behavior occur); the second two for an ultimate explanation (why does the behavior occur). We turn first to research areas that may help evolutionary behavioral economics provide a proximate explanation for economic behavior, and then to those that may help it provide an ultimate explanation. However, such a classification is inevitably somewhat arbitrary, because the crucial point of Tinbergen's argument is that all four kinds of explanation are needed and that they are inevitably intertwined.

Money as a Special Type of Economic Good

To give a preliminary flavor of how Tinbergen's approach works in practice, we consider here how it might apply to one particular economic phenomenon: money. Before starting the analysis, however, we need to ask why we should consider money as something that an evolutionary account of economic behavior needs to take into account. Although all modern economies use and indeed depend upon money, it is quite possible in principle to have an economy that does not. Except for the specific study of financial firms and their products, money tends to disappear from economic analysis; it is just often summarized as medium of exchange, a measure of value, a standard, and a store of value, but not an economic product in itself.

The major reason that an empirically grounded economics needs to deal specifically with money is that, empirically, ordinary consumers and firms do not always behave toward money in the way of conventional economic thinking, or the standard "theory of money" (Friedman 1994; Friedman and Schwartz 1982; Lucas 1972; Mises 1981) presupposes. That does not necessarily mean that people's real behavior toward money could not in principle be captured within a conventional general equilibrium model; it means that if, for other reasons, we are moving to a complexity/evolutionary approach to economic modeling, we need to incorporate accurate empirical information about behavior toward money—or, at least, we need to be in a position to test whether the model is sensitive to the difference between assumed and actual behavior.

So in what sense is behavior toward money anomalous? Lea and Webley (2006) argue that, psychologically or biologically, a major puzzle is that money is too strong an incentive. Behavior toward it resembles behavior toward incentives that have strong and obvious roots in biological needs (whether for the survival of the individual or for reproduction), that go down through our entire vertebrate heritage. It does not resemble our behavior toward objects that are merely means to a biological end. Lea and Webley argue that money is too recent an invention for there to have been any genetic adaptation to its existence. More recent accounts of gene–culture coevolution, which show that the oft-cited evolution of lactose tolerance is just one example of hundreds that operate over relatively short timescales (Laland et al. 2010), give us some cause for pause about that argument. However, we can continue to rely on the

“stylized fact” (deduced from the history of European exploration and colonization of regions where no money was used) that people with no experience of money adapt to it very quickly when a money economy begins to exist in their region, and certainly their children have no problems with it. It would be good, however, to have real historical/anthropological data on how far this is true. Furthermore, there are behaviors toward money that are not explicable on the standard economic account, which is that money is valued and understood solely in terms of what it can buy. Lea and Webley (2006) explored these issues at length and coined the metaphor “money as drug” to capture these various phenomena (though they acknowledge that some of them fit the drug metaphor better than others).

An example of these noninstrumental aspects of behavior toward money, and one that links with other issues raised in this report, is the widespread existence of social norms about the use of money as a gift. The nature of these norms varies between times and cultures, but they are common. Their significance can be seen in the literature on Christmas giving. It can be estimated that there is an annual loss of utility equivalent to around 5 billion USD in the U.S. economy, because people insist on giving specific gifts at Christmas, rather than giving money, which the recipients could spend on whatever they wanted—the so-called deadweight cost of Christmas (Waldfogel 1993, 2002). Of course, we can all think up reasons why this behavior might occur (e.g., Ellingsen and Johannesson 2011), because we live in the kind of society that has this kind of social norm, but that is not the point. Explicably or otherwise, utility is still lost.

So how does human behavior toward money look from the perspective of Tinbergen’s four questions?

1. In terms of mechanism, it is clear that for humans, money functions as a reinforcer in the Skinnerian sense, and that the brain mechanisms underlying its incentive effects cannot be distinguished from those underlying the effects of obviously biological reinforcers (Levy and Glimcher 2012).
2. In terms of development, it is equally clear that behavior toward money is something that is generally learned during childhood (although as noted, adults can be introduced to money), and indeed we have quite a lot of information about how that learning progresses in a modern developed economy (for a brief summary, see Webley et al. 2001:34). Obviously the development of a concept of money depends on the development of cognition more generally, both in regard to number and other mathematical concepts, and also in regard to social life. It also depends on an understanding of the particular society one lives in, so it looks very different in less Westernized societies (see, e.g., Bonn and Webley 2000) than in a fully modern economy.

3. In terms of its phylogenetic history, money as we know it is obviously an exclusively human phenomenon, but some possible antecedents can be seen in nonhuman primates, where physical tokens that look distinctly like coinage can be used extensively as rewards (e.g., Parrish et al. 2013). That is not to say that money can be understood as a token reinforcement in the Skinnerian sense; the realities of other animals' behavior toward tokens rule that account (Lea and Midgley 1989).
4. As for the adaptive function of money, that is the nub of the argument that Lea and Webley (2006) were making about its tool and drug properties. Where money can properly be seen as a tool, its adaptive function is readily understood. Where it has to be seen as a drug, behavior toward it is quite likely to be maladaptive, perhaps representing an example of what, below, we call a "mismatch" between our evolutionary heritage and our current environment.

Regard for Others

Human morality can be expressed through the rational actor model and validated by experimental and behavioral game theory. Human morality is a rational system that involves making personally costly choices to promote ethical goals. People not only balance self-regarding against moral concerns, they also face conflicting moral principles in making choices.

Human actors exhibit three types of motives: self-regarding, other-regarding, and universalist. Self-regarding motives include seeking wealth, consumption, leisure, social reputation, status, esteem, and other markers of personal advantage. Other-regarding motives involve a concern for fairness and a compassionate interest in the well-being of others. Universalist motives are those that are followed for their own sake rather than directly for their effects. Among these universalist goals, which we term character virtues, are honesty, loyalty, courage, trustworthiness, and considerateness. Of course, universalist goals normally have consequences for those with whom one interacts, and for society as a whole. But one undertakes universalist actions for their own sake, beyond any consideration of their effects.

Humans evolved in hunter-gatherer societies, consisting of a dozen families or so (Kelly 1995) in which political life was an intimate part of daily life, involving the sorts of self-regarding, other-regarding, and universalistic motivations described above. In particular, political activity was strongly consequentialist: a single individual could expect to make a difference to the outcome of a deliberation, a conflict, or a collaboration, so that our political morality developed intimately entwined with material interests and everyday consequentialist moral (Boehm 2009).

As we move from small-scale hunter-gatherer societies to modern mass societies with millions of members, the public sphere passes from being intimately embedded in daily life to being a largely detached institutional arena,

Table 8.1 Typology of human action.

	Personal	Social	Universal
Private persona	<i>Homo economicus</i>	<i>Homo socialis</i>	<i>Homo vertus</i>
Public persona	<i>Homo autisticus</i>	<i>Homo parochialis</i>	<i>Homo universalis</i>

governed by complex institutions controlled by a small set of individuals, and over which most members have, at best, formal influence through the ballot box and, at worst, no formal influence whatever. Political activity in modern societies is thus predominately non-consequentialist, meaning that individuals do not base their choices on the effect of their actions on political outcomes. Except for a small minority of individuals contesting for personal power, the political choices of a single citizen affect public sphere outcomes with a probability very close to zero—sufficiently close that these choices cannot be attributed to consequentialist motives, whether self-regarding, other-regarding, or universalist. This gives rise to the following typology of human action (Table 8.1; see also Gintis, this volume):

Homo economicus is the venerable rational selfish maximizer of traditional economic theory. *H. socialis* is the other-regarding agent who cares about fairness, reciprocity, and the well-being of others. *H. moralis* is the Aristotelian bearer of non-consequentialist character virtues. The new types of public *persona* are *H. autisticus*, who behaves politically just as *H. economicus* does privately, while *H. parochialis* votes and engages in collective action on behalf of the narrow interests of the demographic, ethnic, and/or social status groups with which he identifies. Finally, *H. universalis* acts politically to achieve what he considers the best state for the larger society, for instance, reflecting John Rawls's (1971) veil of ignorance, John Harsanyi's (1977) criterion of universality, or John Roemer's (2010) Kantian equilibrium.

Norms

Norms govern social interaction and can easily be discussed within the language of institutions. Crawford and Ostrom offer a simple syntax of institutions containing the elements T, A, E, D, C, where T specifies the target group, A indicates a range of actions, E indicates the target environment, D is a deontic conditional permissive, prohibitive, or necessary, and C is a range of consequences. Institutional statements of the form [T, A, E] are called "shared strategies" which can arise from equilibrium behavior and are self-enforced (Crawford and Ostrom 1995). Statements of the form [T, A, E, D] are called norms and require decentralized enforcement. Statements of the form [T, A, E, D, C] are called rules and require centralized enforcement. Theoretically, these statements and the mechanisms/institutions that support them function to minimize the cost of social interaction.

Two classes of models regarding the effects of norms are used in economics. The first looks at how norms constrain the opportunity set of individuals; the second looks at how norms change the conditional payoffs to individuals thus incentivizing them to align their goal with some external (group or leader's) goal. A third possibility, discussed by Adam Smith in *The Theory of Moral Sentiments*, introduces the "the impartial spectator." An oversimplification is that the impartial spectator serves as an internalized mechanism which tells an individual whether or not a particular (action, outcome) pair is acceptable. The impartial spectator may reduce the number of institutional statements that include a deontic or consequence, but clearly does not eliminate them.

The role of the impartial spectator is a complex one in human decision making. The neuroscience models of human behavior assume that the brain can produce subjective values and then use these values to guide a process of learning the optimal choice. In general, this amounts to the problem of solving a time inconsistent stochastic dynamic programming problem that involves all of the decision-making events, both anticipated and unanticipated, during an individual's lifetime.

Using Bellman's principle of optimality, we could imagine each decision event as being decomposed into an immediate expected payoff of an action and a discounted future expected value of the optimal program given the action was taken. One might suspect the difficult problem is to compute this future value. The role of Smith's impartial spectator may be to select among simplified variations of future expected value functions that maintain the inclusive fitness of the individual. In such a model, norms may be used to assist in this selection process. In the extreme, imagine an individual facing a prohibitive norm over some action choosing to make the expected future value of choosing the action zero. If this reasoning is correct then an interesting research agenda is to understand how shared strategies, norms, and rules affect the form of the expected future value function and when does it lead to more consistent decision making over time.

Social Learning/Cultural Evolution

In the standard model of rational choice discussed above, preferences are assumed to be exogenous. When making sense of the choices of an individual at a given snapshot in time, this is perfectly reasonable, but preferences change across time and space. Cultural evolutionary theory (Richerson and Boyd 2005) represents one way to model preference dynamics while addressing Tinbergen's four questions for understanding behavior. Culture, defined as beliefs and behaviors influenced by social learning, evolved to allow humans to adapt rapidly to their ecology (survival value). The act of learning from others is conceptualized through social learning strategies which "bias" the way individuals learn from others throughout their life (causation and ontogeny). Prestige bias, success bias, frequency-dependent biases, and ethnic marking

are strategies that have been shown to be theoretically salient in the adoption of beliefs and behaviors (Boyd and Richerson 1985; Henrich and Boyd 1998; Henrich and Gil-White 2001; Nakahashi et al. 2012), and that have received support experimentally (Efferson et al. 2008a, b; McElreath et al. 2008; Chudek et al. 2012). Building economic models with preferences informed by social learning strategies can help us think about economies as complex adaptive systems (e.g., Bell 2013). We give an example of how this might work by considering the dynamics of how social norms are formed or affected by migration and ethnic group formation.

As the formation of preferences through social learning occurs within the context of a group, how individuals join and form ethnic groups become obvious determining factors. Notably, some ascribe the formation of ethnic groups, and by extension the preferences formed within them, as having essentially to do with the details of the migration process. Migration rates, migrant age distribution, and the characteristics of the host community appear influential in real cases (Charsley 1974; Wallman 1986; Portes and Rumbaut 2006). A prominent feature of modern migrant groups, and multiethnic groups generally, are external and visible markers of group membership, termed ethnic markers (Barth 1969). There is evidence that ethnic markers are a possible salient feature in premodern groups as well (Hodder 1977). In a laboratory setting, ethnic markers have also been shown to evolve (Efferson et al. 2008a, b), following cultural evolutionary predictions that ethnic marking can evolve under a wide range of coordination benefits (McElreath et al. 2003). These findings corroborate ethnographic accounts that document the formation of ethnic groups without any apparent political benefit to do so (Charsley 1974). Thus, weak benefits to an ethnic identity may be sufficient to maintain separate group identities within a larger population.

Work on the origin of group ascriptions and transnationalism further highlights the flexibility of group membership (e.g., Lee 2009; Moya and Scelza 2015). This suggests preferences to be highly flexible also, mapping onto group membership. Perhaps even more intriguing is the empirical observation that new groups, norms, and hence preferences may be formed through 1+ generation immigrants who mix aspects of the originating and host populations, which is clearly seen linguistically through the generation of mixed languages, pidgins, and creoles (Thomason 2001). This process of novel group formation and possible expansion through group selection (Boyd and Richerson 1990) adds dimensions of study which lead toward a complex systems approach to group membership and preference emergence.

The formation of ethnic markers, adoption of group preferences, and migration itself are all fundamentally shaped by social learning. However, the specific social learning strategies used by individuals (e.g., conformity or prestige bias) are hypotheses and are best used to form competing economic models and tested against data. Finally, there is an added benefit for taking this cultural evolutionary perspective. Given the empirical-theoretical feedback between

cultural evolutionary models and anthropology, cultural evolution provides a natural way to incorporate anthropological insights into economic models.

Learning/Beliefs

Over the course of the last several decades, psychologists, neurobiologists, and other students of animal behavior have made great strides in understanding the mechanisms by which animals learn the values of their actions. From the point of view of economics, these kinds of learning take two forms: learning about one's own preferences and learning (or building) one's beliefs about the environment.

Consider an animal that encounters a new food which it has never before tasted. It samples that food repeatedly and in various quantities, learning by trial and error the desirability of that food. From the point of view of an economist, this can be seen as the animal learning how much it desires that food, what an economist might call preference learning. Consider next an animal learning by trial and error how likely it is to gain a reward by a certain action or learning how likely it is to outcompete another specific animal in some kind of competitive interaction. Both of these are what an economist would call examples of belief learning.

Interestingly, psychologists and neuroscientists do not generally distinguish between these two kinds of learning. They generally bundle them together in their mechanistic study of learning. Perhaps even more interesting is the observation that (at least in animals) the brain areas associated with learning, which have been most heavily studied, does not distinguish between these two classes of learning. (Without doubt, other brain areas do make these distinctions, although little is known about this.)

Studies over the last few decades have, however, led to the development of a simple taxonomy of learning mechanisms. It appears that the mammalian brain employs a number of different discrete complementary mechanisms for learning—whether learning about beliefs or about preferences. Speaking broadly, the most well-studied of the learning mechanisms fall into four categories: model-free reinforcement learning, model-based reinforcement learning, habit learning, and declarative learning (which can be broken into two differentiable categories: episodic and semantic). One important side note is that while all of these mechanisms can influence behavior of the most subtle and deliberative kinds, not all are accessible by verbal inquiry. While subjects often, for example, use model-free reinforcement-learning mechanisms when engaged in repeated strategic games, they often cannot verbally report the values generated by these systems with the same precision that their behavior accesses.

The best studied of these is the reinforcement-learning mechanism (Bush and Mosteller 1955; Wagner and Rescorla 1972), which is described in some detail by Glimcher (this volume) and is often referred to as the model-free learning system (see also Daw 2013; Daw and Tobler 2013). Briefly, this

mechanism employs the neurotransmitter dopamine to recursively compute the value of a given stimulus or action from repeated samples. It can be thought of as developing a time-weighted average of experienced rewards and punishments to derive an expectation. Because it accomplishes this using a recursive algorithm, the single-valued product of the learning process cannot be inverted to separate utility and probability—even if those two elements contribute very differently to the expectation that guides behavior. This system is localized to a set of brain regions in the basal ganglia and frontal cortex. One great limitation of this system is that it can handle only first-order representations of the external world. It is thus exceedingly difficult for this system to represent the values of higher-order contingent structures in the external or behavioral space. Nonetheless, there are many simple situations in which this module can approximate Bayesian learning (Sutton and Barto 1981).

Model-based reinforcement learning, associated with the frontal cortex and hippocampus, appears optimized to learn higher-order conditional structures. For example, learning the value of a two-move sequence, where significant contingent structure relates the two moves to rewards, is the type of problem that this system appears to solve well. In essence, the system seems well-suited to extract a causal model of the environment's reward structure, a process which is computationally much more data intensive than model-free estimates. While the details of the model-based system are just becoming known (Daw et al. 2011), it is clear that, like the simple reinforcement-learning system, it relies on dopamine neurons to report a mismatch between observation and expectation to drive learning.

Interestingly, the behaviorally expressed values which animal and human subjects place on an action with second-order reward contingencies appears to be a weighted sum of the model-free and model-based systems. Early in learning, when little data is available, the model-free system seems to dominate the estimate of value expressed behaviorally. As sampling progresses, this shifts, and the model-based system makes a larger contribution. Neural data shows a similar shift away from the basal ganglia, where the model-free system operates, to the hippocampus, which is involved in the model-based system (Daw 2013; Daw and Tobler 2013).

A third learning system (also resident in the basal ganglia) is known as the habit system. This system is similar in many ways to the model-free system, but it shows very little plasticity once a certain amount of learning is complete. In a typical assay for this system, an animal is trained to produce a behavioral response for hundreds of rounds. During initial training, which is dominated by the model-free system, subjects are responsive to changes in the environment but as this system comes to dominate behavior, the animal becomes unresponsive (or very slowly responsive) to changes in the environment (Dezfouli and Balleine 2012). In the language of reinforcement learning, the effective learning rate drops dramatically as this system comes on line, thus forming a robust habit. Recently, the ability to turn this brain area on and off using invasive

methods has made it possible for experimenters to engage and disengage the habits of their animal subjects at will (Smith and Graybiel 2013).

The final class of learning system is the one that will be most familiar to economists: the declarative-learning system. This hippocampus-based system stores episodic memories; that is, consciously accessible memories of our lives (e.g., the day one's child walked or the time one was first taught the Slutsky equation; Squire and Zola 1996). This system also stores semantic memories; that is, nonautobiographical fact-based memories (e.g., the recall of the Slutsky equation but not where one learned it).

The most important feature of learning for neurobiologists is that it is accomplished by a finite set of discrete modules that execute a specific set of computations. While there are doubtless many of these systems, they are generally tractable to biological isolation and behavioral categorization. These are advantages that empirical approaches to learning may have over more theoretical approaches, which often find complex learning situations intractable.

Motives

Explaining “why a person in a given situation selects one response over another or makes a given response with greater energization or frequency” is the aim of motivation theory (Bargh et al. 2010:286). In economics, the answer usually given to this question is that rational agents choose actions that satisfy their preferences—whatever they might be. By recourse to motivational theories in behavioral science and psychology, it is possible to add substance and to be more specific with respect to the preferences (Witt 2008). According to these theories, some motivations to act are innate and shared across the higher species, including humans. The average strength of these motivations (defined per suitably chosen period of time) is distributed within the population of a species with the usual genetic variance. Some motivations relate to the organisms' homeostasis. Examples in humans of such economically relevant biological motivations are the need for food, clothing, or shelter and the seeking of social status. It can be conjectured that the motivation for seeking social status was positively related to survival and reproductive success in ancestral times, when these motivations seem to have been genetically fixed. Yet another economically significant motivation that seems to be hardwired is the seeking of cognitive or emotional stimulation (or “arousal”).

These motivations—the list could be prolonged—are of interest here, because their impact is observable in national household expenditure, the variance caused by cultural and cognitive influences notwithstanding. In this respect, they do not differ from other innate behavioral traits whose impact becomes visible in the grand total of large populations. A prominent feature of these motivational forces is their satiability in terms of pro rata consumption (i.e., their declining marginal strength, when the ability to spend on them rises with growing income). This feature is expressed by the shape of Engel curves based

on aggregate household data. An example of a very robust satiation effect is the declining share of food expenditures in all household expenditures. It occurs both across income percentiles and over time (U.S. Bureau of Labor Statistics 2006); in absolute terms, spending nonetheless continues to rise slightly.

In contrast, the motivation for seeking social status, likely to be one motivational force behind the strongly increasing share of expenditures on housing, appears to be difficult to satiate as income rises. One reason may be that this motivation is tied to a consumer's relative position, which does not change when all consumers in the reference group similarly increase their spending. As a consequence, a potentially unstable expenditure race can emerge (Frank 2012). The motivation for seeking stimulation is likely to be one motivational force that drives the strongly rising shares of expenditures on entertainment and tourism (for the long-term trend in the mentioned expenditure categories in the United States, see U.S. Bureau of Labor Statistics 2006). Again it appears that little, if any, satiation occurs as the ability to spend increases. In this case, however, the reason seems to be a stupefaction process that requires ever new stimuli to be obtained from consumption so as not to fall below a certain arousal level (Scitovsky 1981).

Culturally contingent parts of the motivation to act are more difficult to identify. They result from individual conditioning learning and cognitive learning. Conditioning learning (of reinforcers) is a hardwired phenotypic adaptation mechanism in many higher species and humans. It contributes to survival and reproduction success. Reflecting the influence of a lifelong individual learning history of human consumers, the motivational forces emerging from this kind of learning are likely to contribute substantially to the huge variety of, and the idiosyncrasies in, the consumers' individual expenditure patterns. Thus far less is known about the satiability of these motivational forces, but behavioral science research seems to point to low or no satiability when income rises (Leslie 1996). Something similar holds for the consequences of cognitive learning and the acquisition of instrumental consumer knowledge. As long as cognitive motives for consumption can be derived from means–ends considerations or exist in their own right (as, e.g., the need for self-esteem), no satiation occurs. Producers' advertising and product information promotion often aim at suggesting convincing means–ends relationships in favor of their offers.

Differences in the relative, income-driven satiation of the consumers' motivation to spend can explain some of the long-term changes in the composition of the household expenditure over periods of strong income growth. Moreover, the existence of both innate motivations and culturally contingent motivations that, for different reasons, seem to be difficult to satiate also explain the secular growth of consumer expenditures (Witt 2011). The motivational mechanisms preventing satiation are essential for understanding the demand-side dynamics for economic growth. They also raise intricate questions.

In market economies, the process of production is ultimately organized to serve the satisfaction of consumer preferences and is therefore said to enhance

welfare. However, because of the specific motivational mechanisms that prevent satiation in the case of the need for seeking social status and the need for stimulation, the corresponding growth of expenditures does not increase the consumers' welfare. In terms of satisfying these specific preferences, consumers tend to just stay put. Still, the costs of growing consumption in the form of resource depletion and environmental degradation continually increase. Further, the culturally contingent motivations emerging from conditioning learning and cognitive learning imply preference changes. More specifically, it is a learning of new preferences by which satiation in the motivation to consume is postponed or offset. When preferences change, a welfare assessment of a growing consumption motivated in this way runs into problems. Its outcome hinges on whether the pre- or the post-learning state of preferences is chosen as reference point. In any case, there will be the environmental costs of a growing consumption.

In the language of evolutionary theory, the hypotheses about the motivations driving the process of economic growth raise the following question: If the economy is a complex adaptive system of type 2 (Wilson, this volume), what is the criterion of the selection process operating on this system and its growth? It does not seem to be the contribution to reproductive success (as it used to be during most of human phylogeny). Nor do the mentioned features of innate and culturally conditioned motivations (i.e., of the agents' preferences) seem to support the idea that the criterion is the contribution to human welfare. Instead, the secular growth of the use of natural resources suggests a criterion that may be conjectured to be valid for all living complex adaptive systems: the contribution to raising the throughput of energy and other resources.

Mechanism

The last fifteen years has seen a revolution in our understanding of the mechanism of human and animal choice behavior. Just a decade ago, even the most basic outlines of the choice architecture were unknown. Today, neurobiological correlates of utility, the maximization underlying the actual making of a choice, and even some of the neural mechanisms underlying the anomalous choice behaviors identified by behavioral economists are beginning to be understood mechanistically. Even more exciting is the fact that the combination of normative theories of information encoding from neuroscience and normative theory from economics are yielding new insights at a predictive/positive level that may come to be important to welfare economics (Glimcher 2015). Here we review some of the key insights that have emerged over the last decade from the new discipline of neuroeconomics (for further details, see Glimcher, this volume and 2011a).

Since Samuelson's pioneering work in the early part of this century (Samuelson 1947), economists have known that consistent human choosers behave exactly "as if" somewhere in their brains, stable values are assigned

to any option that lies within a (feasible) choice set, and that these subjects behave exactly “as if” choice was the process of selecting the element (from this choice set) that has the highest numerical value (Friedman and Savage 1948; Mas-Colell et al. 1995). While real human choice is at times inconsistent (e.g., Allais 1953), a group of neuroeconomists began their work by searching for the physical instantiation of utility, and for a mechanism that can select from a physical representation of option utilities the option that has the highest numerical value.

Real success was achieved about five years ago, when a pair of brain areas were identified which had activity patterns linearly correlated with behaviorally inferred utility (Levy and Glimcher 2012; Bartra et al. 2013). In these two areas, brain activity (which is a physically measurable quantity, and hence a unique and fully cardinal object), is now often referred to as encoding “subjective value.” Subjective value (as represented in the medial prefrontal cortex and ventral striatum) can be used to predict choice (Levy and Glimcher 2012). In fact, when humans behave in an inconsistent manner (e.g., because of local context effects), activity in these brain areas still predicts choice accurately (Kable and Glimcher 2007).

At a more biophysical level, studies have started to explain the process of selection among alternatives (Webb et al. 2014). Models of interneuron competition constructed at the level of a couple nonlinear differential equations have begun to reveal the dynamics of this process (LoFaro et al. 2014) and to give insight into how choice operates mechanistically.

Of course if choice begins mechanistically with a representation of subjective value, it is critical to understand how these subjective values are learned or aggregated. A great deal of success has been achieved in studying the neural mechanisms of learning; at least half a dozen such mechanisms have now been identified and at least partially characterized (see Daw 2013; Daw and Tobler 2013). The key point is that mechanisms exist by which the desirability of actions and rewards can be sampled from the environment and used to establish or modify subjective values.

A final critical point to make in this domain is that one should not draw the conclusion that the architecture for choice is monolithic. At each stage of the process, multiple mechanisms cooperate to yield each component in the choice process. For example, many different brain areas, each with unique specializations, combine to set subjective values (Bartra et al. 2013). Multiple interlocked systems represent those aggregated subjective values, and several mechanisms appear to converge the selection process (Glimcher 2015; Padoa-Schioppa and Assad 2006), also in an interlocked manner. For these reasons, it is important not to think of any stage of the process as reflecting the preferences of a single module or actor. It is more accurate to think of the individual human as aggregating specialized components—an aggregation that may soon begin to explain some of the inconsistencies we, as a culture of scholars, see in human behavior.

Modularity

While there has historically been variation in the use of the term, in some communities “modularity” has come to refer to the extent to which a mechanism is functionally specialized (Barrett and Kurzban 2006). It was introduced into broad use within cognitive science by Fodor (1983), who concluded that while some systems in the mind (e.g., sensory systems) were modular, according to his use of the term, others were not. This conclusion rested on the idea that the extent to which a given system was modular depended on how many properties, which Fodor associated with modularity (e.g., automaticity, fast operation), the system in question had.

In current use, the issue is less about how many of the properties that Fodor associated with modularity a given system possesses, and more of a guide for investigation. Present conceptions of modularity focus attention on the empirical question how modular, or functionally specialized, a putative mechanism is. This approach follows the suggestion that modularity was not an all-or-none property, but a property that a computational system can have to a greater or lesser degree (Fodor 1983).

To take one example from a well-understood model, consider the visual system. The front end of the visual system consists of photoreceptors, which are sensitive to the presence of light, and fire depending on the wavelength of light that hits them. Their function is to detect light and send information about its presence downstream in the visual system. Photoreceptors are modular in the sense that they are specialized to detect light and signal its presence. The visual system, more generally, is specialized as well, designed to use incoming light, as well as knowledge in the rest of the brain, to construct an image of the outside world, which can then be used to identify objects, plan motion, and so forth.

A potentially important feature of modular systems is that they can be “walled off” from other systems. (Fodor used the term “informationally encapsulated” to refer to this property.) To return to the example of the visual system, consider the Müller-Lyer optical illusion (Figure 8.1), to which many, but not all, people are susceptible (for further details, see Henrich et al. 2010). If one disregards the arrowheads or fins at the end of these lines, are the lines the same length or is one longer?

The two horizontal lines in Müller-Lyer image are the same length. Even after subjects are told that the two lines in the image are identical, and allowed




Figure 8.1 The Müller-Lyer illusion.

to verify this by measurement, the lines continue to be perceived as unequal. Knowledge of their equality does not affect that visual system; the visual representation is “walled off” from the relevant propositional knowledge.

Some have argued that most of the systems in the mind are modular to an interesting extent (Kurzban 2010). In the context of our discussion here, this might have significant implications for decision making, including in economic contexts. For example, the modular view suggests that different modules are dynamically activated and deactivated depending on context and state. For instance, when organisms are hungry, their “food-seeking” modules might be more likely to motivate food-seeking behavior, simultaneously inhibiting other kinds of motives. In this way, modular systems can be thought of as animating the trade-offs that organisms must make in decision making.

A second implication of modularity resonates with “multiple selves” models (e.g., Ainslie 1992) under development. If the mind is modular and there are a large number of specialized systems operating simultaneously, to some extent “walled off” from one another, then people might not always be profitably conceptualized as unitary, with beliefs and preferences that can be thought of as associated with the person as a whole. In that case, it might be sensible to consider that different parts of the brain might contain different (possibly conflicting) beliefs. Such an architecture might not resonate with models that understand agents, as a whole, to have “beliefs” and “preferences.”

More generally, an implication of modularity is that there might be inconsistencies in the mind. As illustrated by the Müller-Lyer illusion, there might be other domains in which there are contradictory beliefs (representations) in the mind. For instance, one way to understand “self-deception” is that there is one representation that is (more or less) the “true” belief, and one that people communicate to others (Trivers 2000). This view implies that there is no one true unitary agent in the mind. Such a view diverges from some conceptions of economic agents and could have implications for thinking about decision making.

Mismatch

Evolutionary mismatch exists when an organism is “out of sync” with its environment in some significant fashion. Because behavior is produced by specific mechanisms, mismatch can create puzzling, self-destructive, or otherwise aberrant behavior in particular environments (Bowlby 1969; Burnham 2013; Irons 1998; Tooby and Cosmides 1992). It is difficult to think about mismatch except from an evolutionary perspective.

Mismatch can be viewed using Tinbergen’s four questions (discussed earlier). In particular, mismatch and Tinbergen may be particularly important for explaining non-maximizing or other puzzling behavior. Survival value, one of Tinbergen’s four questions, argues that natural selection favors optimizing behavior. For example, in the case of diet (explored below), survival value argues that organisms will eat foods that are good for them. However, the selective

pressure to behave optimally is constrained by Tinbergen's other three questions. Mechanisms to produce behavior evolve over evolutionary time and are constrained by ontogeny and phylogeny.

A central concept is the idea of an ancestral environment, which may have varied in many ways, but where important relationships remained relatively constant for many generations—enough for evolution to move the population toward local maxima.

The developmental psychologist John Bowlby coined the term *environment of evolutionary adaptedness* as “the one that man inhabited for two million years until changes of the past few thousand years led to the extraordinary variety of habitats he occupies today....Just as Darwin found it impossible to understand the structure of an orchid flower until he knew what insects flourished and visited it in its environment of adaptedness...we need to turn to anthropological studies of human communities...and field studies of the higher primates” (Bowlby 1969:59, 61).

Organisms that exist in novel environments can suffer or exhibit anomalous behavior, because their physiologic mechanisms produce different outcomes in the current environment than they did in the ancestral environment. In humans, one example of mismatch is cretinism, a condition of severely stunted mental and physical growth caused by hypothyroidism. Some populations in the Himalayas have cretinism rates of over 10% due to iodine-deficient hypothyroidism. This is a poignant example of the cost of mismatch. (Iodine injections prevent cretinism in this population as readily as limes prevent scurvy.) The mismatch explanation for cretinism is that the human ancestral world people absorbed sufficient iodine without any behavior directed toward its acquisition. In many parts of the world, the baseline iodine levels are high enough that sufficient quantities are absorbed as a part of other activities. Cretinism occurs when humans move to geographic areas with particularly low levels of iodine (Delange 1994; Ibbertson et al. 1972). The high prevalence of goiter in Derbyshire, U.K. is sometimes attributed to mismatch, although this case is not settled (Saikat et al. 2004).

Another example is the idea that people's dietary tastes produce morbidity and mortality. This was articulated by William Irons (1998:197):

It is hypothesized that in ancestral environments these preferences motivated people to come as close as their circumstances allowed to optimal diets. However in modern environments, the abundance of different types of foods is vastly different, and these preferences often motivate people to choose diets that are much less healthy than are possible in their circumstances.

To further elucidate this thinking, consider how economics addresses preferences toward meat, bread, and high-fructose sugar. In standard economic models, these three foods are treated equivalently; people have preferences and make choices consistent with those preferences. An evolutionary perspective suggests that the time depth of a type of food, or, more generally, an economic

good, might be part of the evaluation. Humans are argued to have been meat eaters for many thousands of generations, to have had significant access to carbohydrates such as bread since the invention of agriculture, and to have had much less access to high-fructose sugar (Lucock et al. 2014). The evolutionary suggestion is that human genes and cultures are less likely to have had time to adjust to bread and high-fructose sugar than to meat (for meat, however, it is likely that ancestral meat was leaner than that which is available today).

What are the implications of mismatch for economics? Our preferences can lead us to behaviors that are suboptimal; we do what we like and the outcome is “bad.” Bad can have three overlapping meanings. First, the outcome can lead to disease and death, or it can simply be less desired than another feasible outcome. Second, our preferences can result in behaviors that lead to internal conflict. For example, it has been suggested that humans are natural savers via body fat. If this is true, then human preferences over financial savings instruments might lead humans to save systematically too much or too little, where “too much” or “too little” is defined from the perspective of other parts of the brain (Burnham and Phelan 2000). Overspending, for example, might cause internal conflict. Third, our mismatched preferences can impose costs on others (e.g., drug addiction). These costs can be addressed via conventional methods to align public and private payoffs (e.g. education, sanctions, taxes). However, a mismatch view can suggest likely areas for policy intervention. For example, trans fats are argued to be a particularly problematic type of food as they possess a dangerous set of attributes:

- Trans fats are not distinguished by most humans from other fats.
- Trans fats are cheaper to use.
- Trans fats are bad for health.

Thus, without some form of policy intervention, people might be likely to overconsume trans fats from both their own perspective and that of the society. In addition, those who produce foods containing trans fats have a strong incentive to encourage their consumption.

The mismatch perspective suggests that evolutionary history has importance both for economic theory and for public policy.

Objective Function

Given a defined ecology, an organism is assumed to make decisions that result in an inclusive fitness measure. It is often convenient to provide a proximate problem of choosing a strategy to maximize (or minimize) an objective function subject to resource constraints. Strategies, in turn, are mappings from information sets into actions. The action chosen is often called the decision of the organism. Decisions produce outcomes for the organism which may increase or decrease its inclusive fitness. Selection pressures may change the objective function that is maximized, the strategy space available to the organism, the

mechanisms for choosing strategies, and the mechanisms for converting strategies into actions. Sometimes a shortcut is used to look at selection as acting on the choice of strategy used by the organism.

Several extensions are made. First, to what extent does the objective function and strategy space anticipate the possibility that the ecology may change over the organism's lifetime or adaptive cycle? This may include systematic or random changes in resources. Second, to what extent does the organism have to anticipate the strategies/actions of other organisms? This results in the choice of competitive or cooperative strategies and the formation of beliefs (naïve or sophisticated) about others. Third, to what extent does the organism plan? A plan is a strategy that chooses a contingent path of decisions. Here, the organism tries to solve a stochastic dynamic programming problem to determine its strategy. In solving this problem the organism must form beliefs (either naïve or sophisticated) about future states of the world.

The Origins of Choice

The many cases of mismatch show that simply identifying behaviors (or outcomes) as adaptive (or not) is insufficient to predict whether they will be chosen. Only if there have been appropriate selective pressures operating during the organism's evolutionary history is it likely that there will be a mechanism for making the choice that will enhance the organism's fitness; if those pressures are no longer operating, the choices made may no longer be fitness enhancing. In addition, all of this assumes that the animal has the perceptual and cognitive apparatus to make the choices in which we are interested.

While fossils rarely, if ever, reveal the preferences or choices of our ancestors, we can make estimates of how old particular tendencies are by investigating whether they are found in our near relatives. If the human taste for the sweetest substances available is shared by all, or nearly all, Old World primates, this preference is probably at least 25 million years old (and could be older). If it is shared only by the great apes, the preference is probably no more than 20 million years old (and could even be more recent). It is a reasonable guess that the older a preference, the harder it will be to overcome.

The second issue raised by our evolutionary history is the kinds of learning mechanism available. As noted above, there are at least four distinguishable kinds of learning (the literature of comparative cognition would recognize many others, or divide our four categories into many subsets). Consider, for example, the special kinds of learning shown by a songbird or a long-distance migrant, or the social learning that enables capuchins to acquire tool use or meerkats to deal safely with scorpions. We want to focus, however, on two particular categories of learning, because the distinction between them has important implications for what we might put in a complex theory of human economic behavior.

As far as is known, what we have termed “model-free” reinforcement learning seems to be shared with all vertebrates, and it has been shown to be shared with some invertebrates, many insects (e.g., ants, bees of various species, cockroaches), other arthropods (e.g., spiders), some mollusks (including slugs and snails), and various cephalopods. The most contested question may be whether it extends to humans. Associative learning (which includes both model-free and model-based reinforcement learning) can certainly be demonstrated in humans, but distinguishing model-free and model-based learning by behavioral means alone is tricky. (For a summary of early behavioral work, which demonstrated the necessity to include a mental model in the explanation of rats’ associative learning, see Dickinson 1981.)

The reason that associative learning is difficult to study in humans is the pervasive human tendency to use a different learning mechanism to solve choice problems, namely declarative learning or rule-governed behavior (Skinner 1969, chapter 6). In the majority of cases, this dominates any associative processes (e.g., McLaren et al. 2014). The rules or propositions people use to control their choices may be of many types and may come from many sources: from a personal calculation to a social norm. A particularly important type, however, are those which derive from what Tulving (2001) called “autonoëtic consciousness”: the ability to play out, in one’s imagination, past or future events so as to, test possible responses to a situation, predict their consequences, and use those predictions to guide a decision. When autonoëtic consciousness is applied to the past, we call it episodic memory; when applied to the future, it is called episodic future thought. Suddendorf and Corballis (1997) refer to these processes as “mental time travel,” and argue strongly that they are unique to humans—an argument that experimentalists have been assaulting ever since, as we now explain.

It is fairly well established that a range of animal species can show what is generally referred to as “episodic-like” memory, demonstrating a recall of what happened, when it happened, and where it happened. This is most easily demonstrated in animals specialized for scatter-hoarding: the earliest evidence came from scrub jays (Clayton and Dickinson 1998), and was followed by demonstrations in chickadees (Feeney et al. 2009). This has also been demonstrated in other species using different models: in mate selection by meadow voles (Ferkin et al. 2008), spatial food search by rats (Babb and Crystal 2005), and matching to sample in rhesus monkeys (Hoffman et al. 2009). What remains unclear is whether these episodic-like memory phenomena represent the kind of declarative process that would enable animals other than humans to engage in rule-governed decision making. Direct comparisons between humans and animals, like pigeons, in tasks where human choices are typically rule-governed reveal unsurprisingly massive species differences. For example, gray squirrels, which as scatter-hoarders might be expected to have episodic-like memory available, reportedly behaved more like pigeons rather than humans (Wills et al. 2009).

From the point of view of animal choice, perhaps the most important question is whether animals other than humans can imagine the outcomes of a choice and behave in the light of those imagined outcomes. Currently there is a lively research area on planning in animals that focuses on nonhuman primates, with clear claims from laboratory experiments that planning can be demonstrated in chimpanzees and rhesus monkeys (Beran et al. 2004; Mulcahy and Call 2006) as well as in capuchins (Beran and Parrish 2012). This distribution suggests that a considerable number of primate species should show the ability. Field data on route selection while foraging have also been interpreted as indicating planning in, for example, chimpanzees (Ban et al. 2014), although apparently not to any significant extent in baboons (Noser and Byrne 2010). Planning, however, may not be available for all tasks, even in chimpanzees (Dufour and Sterck 2008).

While the evidence for planning in animals is not yet fully robust, it is strongly suggestive. In the context of this Forum, what this means is that declarative learning and rule-governed decision making may be phylogenetically quite ancient, and specifically more ancient than human language. It is beyond question that a complex model of human economic choice, and therefore of the economy, needs to include the possibility of rule-governed decision making. If, however, rule-governed behavior can be found in nonhuman animals, as the data now suggest, it follows that such behavior can occur without the intervention of language. Therefore, in humans, it is possible that there may be kinds of rule-governed behavior that are resistant to linguistic probing, or indeed to alteration by linguistic argument.

Conation

One of the most striking departures of human choice from rational models occurs in intertemporal choice (i.e., choice between outcomes that occur at different points in future time). Humans are much more present-oriented than seems sensible, and to make observed choices fit with a rational choice model, absurdly high discounting rates have to be postulated (e.g., 33% per annum; Friedman 1957). This excessive present orientation is frequently attributed to “weakness of will” (a phrase particularly used in philosophical psychology) or “lack of self-control.” It has been argued (e.g., Ainslie 1975) that it arises because future outcomes are discounted according to a hyperbolic function rather than to the exponential function required for standard rationality approaches. Reconciling hyperbolic discounting with any version of the rational actor model is difficult. The hyperbolic discounting model derives ultimately from experiments with rats and pigeons, which aimed to produce a nonrational positive theory of choice (e.g., Herrnstein 1970; Mazur 1984). When we compare human choice with these kinds of animal choice, it turns out that the question is not why humans are so bad at intertemporal choice, but rather why we are so (relatively) good at it (Lea 2014). Furthermore, the most recent experimental

economic evidence suggests that exponential discounting may fit large-scale data quite well, and with more reasonable discount rates than had been previously thought (on the order of 9%; Andersen et al. 2014). However, it remains true, as Loewenstein and Thaler (1989) and many other behavioral economists argue, that an empirically based model of the human economic agent needs to incorporate both discount rates and discount functions that would not be predicted by the rational actor theory.

Toward an Evolutionary Behavioral Economics

Implications for Economic Theory

So what kind of modification in economic thinking do we need? There are two possibilities, a gradualist and a radical approach, and there were members of the group who argued for each.

Under the gradualist approach, work would go forward using the standard or rational actor model, informed, however, by what we know about the way humans make decisions to modify the inputs and parameters that the model requires. In broad terms, this is how economics, as a discipline, has responded to the challenge from behavioral economics, at least since Kahneman and Tversky's (1979) original paper on prospect theory. As a result of this work, the status of the rational actor model has been clarified. Behavioral economics has repeatedly shown over the last few decades that individuals act inconsistently, make choices that incur costs on them but may benefit others, use rules of thumb and other heuristics, and are prone to a range of biases. Both economists and psychologists who have thought about the theoretical foundations of such behavior have, in most cases, responded by modifying the existing framework of utility maximization, for example by introducing additional elements into a utility function (e.g., identity concerns, social norms, inequality aversion) or adding additional parameters to the utility function (e.g., the "beta" of hyperbolic discounting). However, in such cases, it is not always clear whether the proposed utility functions are thought of as reasonable descriptive approximations of observed behavioral phenomena (*as-if* propositions in the Friedman sense), or whether those utility functions are derived from observed consistencies of human choice behavior (*as-if* propositions in the sense that consistent observed choices are equivalent to preferences). Given that most economists subscribe to the view that economics should be based on behavioral and observable data, the second *as-if* sense should be preferred, but bizarrely, very often proposed models do not live up to this view. Hence, very generally speaking, one way forward would be to clarify such methodological issues and to see how far consistent behavior should be or can be the basis of economic models in the light of empirical and behavioral evidence of inconsistent behavior. Would it suffice to weaken consistency axioms to allow for a broader

range of behavior (e.g., Arlegi and Teschl 2015)? Related to this is also the issue of whether it is desirable to maintain the idea of maximization or whether other choice rules should be considered as well (e.g., Baigent and Gaertner 1996; Gaertner and Xu 1999). What is clear, however, is that an evolutionary approach to economics leads to many different emphases from a conventional one, as is shown in Table 5.1 (see Axtell et al., this volume). If we are to use the rational actor model in an evolutionary context, we need to honor as many of those differences as we can.

The more radical alternative is to discard entirely the standard or rational actor model central to existing economic approaches to understanding human nature. It is no good advocating such an approach without specifying what might be put in its place. One way to formulate a new model, advocated by some members of our group, is to build empirically derived multi-agent models without any of notion of rational actors or maximization. It is perfectly possible to construct multi-agent models of specific economic behaviors without explicitly imposing a rational actor model. This has been done by Gintis (2007), a paper that has served as a model for many multi-agent models of real systems (e.g., model of economic convergence within the European Union; Dawid et al. 2014). To represent the full spectrum of economic behavior, this type of model requires a great deal of knowledge about the agents (e.g., their motives, their stored information about the world, their learning mechanisms). It also requires us to know how agents relate to one another (e.g., their social interactions, the different levels of power and influence they are able to wield). Here the voluminous literature on social network analysis should be helpful (Burt et al. 2013). To put it another way, theory-building would have to depend on a detailed knowledge of the natural history of economic behavior.

In this context, we note that it is normal in economics and game theory to treat the ontology of decision making as utilitarian. Other fields, however, investigate a wide variety of ontologies of human decisions. One such ontology involves deontic logic, which attempts to formalize notions of what is permissible, omissible, impermissible, obligatory, and optional, among other things. Deontic logic is a kind of modal logic and can be axiomatized using standard modal logic. Specifically, using suitable Kripke structures, a logic known as KD45 has come to be known as standard deontic logic and has been widely used, for example, in multi-agent systems, from computer science to animate agents. This logic seems especially powerful for modeling certain kinds of ideologically motivated, adversarial agents, who take actions predominantly on the basis of what they perceive as the “right” thing to do. Concepts such as norms, conventions, and institutions, which tend to be ephemeral and contingent in a utilitarian ontology, arise naturally in a deontic ontology, and it is partially for this reason that they have found use in models where the formation and evolution of norms is important.

Whichever approach one may want to take (and our two proposed approaches are certainly not all there is to understanding and modeling human

behavior), we would like to repeat the cautionary note we raised earlier: As a species, we face urgent concerns, and perhaps even an emergency.

One of the most obvious of evolutionary facts is that the human niche has expanded over time. The human imprint on the planet—as measured by population, biomass, carbon dioxide emissions, or simply GDP—has expanded enormously since the invention of agriculture.

If one is concerned about the human imprint on the planet, and in particular, a need to alter the trajectory of that imprint, then is it reasonable to ask what the rational actor model has to contribute to addressing this problem. The short answer is: very little. By analogy to standard price theory, economics has a usable theory of how people respond to price or income changes, but it does not provide a deep understanding of the drivers of those responses. In short, economics classifies the behavior, but does not explain it. We believe it is important to explain why the economy is growing, and how we can change its trajectory before something “bad” happens (e.g., as a result of overpopulation, more extreme global warming). This requires us to understand the nature of tastes, which the rational actor model regards as beyond investigation. We can use evolutionary theory to inform economics about how humans have increasingly come to dominate our environment. It is the desire to understand the origin and nature of tastes that leads to our investigation into motivation. We feel it is insufficient to agnostically catalog the types of human consumption. Success in understanding, and even survival, may depend on our ability to understand, at a much finer level, what desires satiate, and at what levels.

Progress on these fronts may allow economics to achieve a reconciliation with environmentalism. Environmentalists often make normative arguments; at present, they often do so without engaging actual evolved human nature, but that is not inherent in their approach (Lea 2001). Given the current state of economics, environmentalists may have no alternative to such approaches.

One important conceptual topic is to understand where selection is working. In the framework of Wilson (this volume), is human society a CAS1 or CAS2 system? Is there any reason to believe that the overall system is subject to selection, or should we view individuals or groups of individuals, but not the entire population of the planet, as the unit of selection? Beyond categorizing human society as CAS1 or CAS2, understanding the processes of selection might offer a novel and better way of understanding human motivation. Selection takes place at multiple levels. One of these levels is the group. Within groups, individuals have to follow the “rules of the game.”

Maynard Smith and Szathmáry (2000) classify evolution according to major transitions that have occurred in the encoding of information. Such encoding allows selection to operate on strategies and mechanisms. They argue that the last transition for humans is evidenced by the emergence of language and symbolic manipulation. This created a capacity for forming, and possibly adhering to, institutional rules.

Humans are unique in their capacity to create and follow the rules of a game and thus can be characterized as *H. ludens* (Gintis, this volume). The capacity to follow rules is essential for the creation of institutions. Through cultural evolution, some institutions are selected while others are not. What can this selection of institutions tell us about human motivation and the change of motivation over time? The human ability to manipulate the rules of games presumably plays a central role in the explosion of human population and the consequent impact of our species on the planet. Is it possible to alter the rules in some way to improve our current trajectory?

Implications for Behavioral Public Policy

The influence of behavioral economists on public policy has increased over the last decade. Based on the insight that choices are often influenced by the context in which decisions take place (Kahneman 2011), policy makers have begun to “nudge” individuals to make certain decisions (Thaler and Sunstein 2008). Nudges are defined as intentional changes of the “choice architecture” that affect the behavior of humans; they do not, however, affect the behavior of rationally deciding *H. economicus*. For example, default rules have at most a small influence on the behavior of rational individuals, but a strong effect on the behavior of humans in the domains of energy consumption and organ donation (Johnson and Goldstein 2003; Sunstein 2014). Also, providing individuals with information about what other people do (i.e., communicating social norms) can have large effects on individual choices (Schultz et al. 2007).

The analysis of behavioral, nudge-like interventions has focused on the immediate effects on the nudged individuals. Other aspects of nudges have received limited attention. For example, we often do not know much about the mechanisms that explain why nudges are effective, nor whether applied nudgers are interested in these mechanisms. We have very little knowledge about how behavioral interventions change behavior over time. We know nothing about unintended consequences of behavioral interventions, which might unfold at higher levels of aggregation, affecting the adaptiveness of groups rather than individuals. The evolutionary approach that we suggest in this paper can contribute to a better understanding of the opportunities and dangers of behavioral interventions. Below we list three areas where evolutionary behavioral economics can be of value for behavioral policy making.

First, evolutionary insights can improve the effectiveness of behavioral interventions. The analysis of evolutionary mismatch (discussed above) can provide ultimate explanations of the mechanisms that explain why nudges are particularly effective in certain domains. For example, nudges to encourage a healthy diet might be designed more effectively when acknowledging that human genes had less time to adjust to high-fructose sugar than to bread and meat. Nature might not have had the time to develop a healthy satiation dynamic for high-fructose sugar (see section on Motives). The modular view

(as described in the section on Modularity), suggests that nudges can change decision contexts to activate certain modules in the brain that interact with inner states (e.g., hunger). A nudge might activate a module to think of future consequences (call it “mental time travel” or “autonoëtic consciousness”; see section on Phylogeny) when organisms are hungry to avoid unhealthy eating. More generally, a better understanding of human motivations and how they differ from each other (see sections on Motives and, for money as a particularly motivating reward, Money) can help predict in which domains nudges will be successful (Lades 2014).

Second, the evolutionary economic focus on understanding the changes of tastes, preferences, and beliefs can help us better understand long-run effects of behavioral interventions. Since actions are not only driven by preferences, but also create or change preferences (Ariely and Norton 2008), nudges can affect preference-learning dynamics. While the behavioral one-shot perspective neglects these dynamics, the evolutionary approach advocated in this chapter focuses on long-term changes of what individuals want “in general.” This long-term change might be much more important than short-term fluctuations of what individuals want when “influenced by the context.”

Whether and how nudges affect preference-learning dynamics depends on the learning mechanism at play (see sections on Learning as well as Phylogeny). Since nudges work best in the absence of awareness (Bovens 2013), they might influence preference dynamics via associative-learning mechanisms that are not necessarily accessible by verbal inquiry. Whether nudges can also influence declarative learning has yet to be established. Nudges can also affect social-learning mechanisms as described in cultural evolutionary theory (see section on Social Learning/Cultural Evolution). Cultural evolutionary theory provides a tool to analyze the long-term effects of social nudges (e.g., social norm interventions) on preference dynamics via biased-learning strategies in groups.

Understanding human motivations and their different degrees of satiability can also be useful for predicting long-run effects of nudges: Nudges that relate to motivations which are difficult to be satiated (e.g., status seeking, seeking stimulation and arousal) might induce different learning dynamics than nudges which encourage behavior that satisfies satiable motivations. In particular, the role of nudges to foster or impede our seemingly insatiable desire for money warrants a dynamic analysis.

Finally, multilevel selection models combined with complexity theory might inform us about the effects of changed individual behavior for the success of groups. Group success might feed back to the individual level and change individual behavior. Nudges in the spirit of libertarian paternalism (Thaler and Sunstein 2008) try to improve individual decisions to make individuals better off, and thus promote individual adaptiveness. Nudges that are not in the spirit of libertarian paternalism try to encourage behavior that is costly for the individual but beneficial for the group (Sunstein 2014). An evolutionary multilevel

selection model could provide some information about which type of nudge makes individuals better off in the long run.

It is crucial to acknowledge the normative aspects of an evolutionary behavioral economics of behavioral interventions. Nudges are often criticized for being paternalistic and manipulative (Sunstein 2014). The evolutionary approach generates new arguments in this debate. The discussion of modularity, for example, suggests that there might not be a true, unitary agent in the mind (Lea and Webley 2005). This generates a problem for libertarian paternalism and its aim to encourage decisions that make individuals better off, as *judged by themselves* (Thaler and Sunstein 2008:5). If there are multiple agents in the mind, which one determines what makes the individual better off? Also, if behavioral interventions aim to change preference-learning dynamics, the normative discussion about the “ethics of nudge” has to be expanded to include long-term changes of tastes, preferences, and beliefs (Binder and Lades 2015). Discussion is needed as to whether nudges that affect preference-learning paths are normatively acceptable. Finally, it is important to discuss whether policy interventions which encourage group beneficial behavior, but are costly for the individual, are normatively acceptable. More generally, should “nudging” be part of public policy making at all. There is a similar debate about whether governments should take care of people’s “happiness” and conceive policies to enhance people’s well-being in terms of happiness. Many scholars actually argue against it, with reference to abuses of power and incentives to manipulate any happiness index to the governments’ own advantage (Frey and Gallus 2013). Similar criticism can be applied to governments’ nudging people to do the right or better thing, especially in cases where there may not be a univocally best option to choose or to do.

Clearly, people should not be deprived of their own decision power and understanding of themselves. Economics has always cherished the idea of consumer sovereignty, and it is the view of many economists and other scientists that this value should be maintained. The alternative to nudging is individual responsibility (Sugden 2004, 2007). As Sugden shows, if responsibility is an accepted normative disposition, then what matters in determining whether a situation is “better” than another are more opportunities, even the opportunities to make bad choices. Clearly, this opportunity–responsibility approach needs to be evaluated from an evolutionary and complexity perspective to see in how far it would fare better than alternative policies in succeeding to provide and maintain a sustainable way of life. Another perspective of public policy is given by Colander and Kupers (2014). As a viable solution for policy making, they propose a combination of government and market mechanisms.

Whatever the approach to public policy making, it seems clear that people have to become aware that to a large extent, their behavior is on many occasions “out of sync” with the kind of behavior needed to save and protect the natural environment, and thus their own lives. Not only should the current model of choice and economic behavior be revised to take account of the huge

amount of experimental and behavioral evidence that people do not behave as the standard economic model presupposes, responsible agents (whether they act as consumers, producers, service providers, government agents, activists, etc.) need to endorse Tinbergen's four questions and be inquisitive about how behavior occurs and why it occurs as it does—and to use this knowledge to change purposefully any future behavior.