האוניברסיטה העברית בירושלים

THE HEBREW UNIVERSITY OF JERUSALEM

RULE-RATIONALITY AND THE EVOLUTIONARY FOUNDATIONS OF HYPERBOLIC DISCOUNTING

by

HILLEL BAVLI

Discussion Paper # 513 June 2009

מרכז לחקר הרציונליות

CENTER FOR THE STUDY OF RATIONALITY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: http://www.ratio.huji.ac.il/

Rule-Rationality and the Evolutionary Foundations of Hyperbolic Discounting

Hillel Bavli*

Abstract

Recent studies involving intertemporal choice have prompted many economists to abandon the classical exponential discount utility function in favor of one characterized by hyperbolic discounting. Hyperbolic discounting, however, implies a reversal of preferences over time that is often described as dynamically inconsistent and ultimately irrational. We analyze hyperbolic discounting and its characteristic preference reversal in the context of rule-rationality, an evolutionary approach to rationality that proposes that people do not maximize utility in each of their acts; rather, they adopt rules of behavior that maximize utility in the aggregate, over all decisions to which an adopted rule applies. In this sense, people maximize over rules rather than acts. Rule-rationality provides a framework through which we may examine the rational basis for hyperbolic discounting in fundamental terms, and in terms of its evolutionary foundations. We conclude that although aspects of hyperbolic discounting may contain a certain destructive potential, it is likely that its evolutionary foundations are sound—and its application may well be as justified and rational today as it was for our foraging ancestors.

I. Introduction

Theoretical economics, and game theory in particular, is under siege. The most fundamental assumption of economics, rationality—the assumption that people behave in a manner as to maximize their utility, given their information—has been challenged, and to some degree turned on its head, by the emergence of behavioral economics. The literature is rife with experimental results indicating an array of behavioral "quirks," and associated claims that people cannot be described as more than loosely rational. Indeed, the idea of "bounded rationality" has stormed the literature and pervaded economic models that range the gamut.

* Center for the Study of Rationality, The Hebrew University of Jerusalem, Fulbright Fellow. I would like to thank Professor Robert J. Aumann, who provided me with advice and encouragement throughout the writing of this paper. I would also like to thank the Fulbright Program for its support.

At the forefront is *hyperbolic discounting*, a behavioral pattern whereby individuals discount future reward with a rate that is not constant as classically presumed, but increasing with the dwindling of time before such reward is realized. It is an empirical finding that emphasizes the contrast between people's impatience for immediate reward and patience for reward when the reward is to be gained at some point in the future.

For example, when offered the choice far in advance, most would choose to spend fewer grueling hours arranging taxes on April 1st than greater such hours on April 15th. As April 1st draws closer, however, one is prone to changing his mind and pushing off the task to April 15th, notwithstanding the greater burden that will result (O'Donoghue and Rabin, 1999). Similarly, experimental subjects display a preference of $10 today over $11 tomorrow, but a preference of $11 in 31 days over $10 in 30 days.

Indeed, the marked implication of hyperbolic discounting is a reversal of preferences. A decision-maker will therefore often find himself in a competition of sorts with his future selves. This form of discounting has been widely characterized as "dynamically inconsistent," and ultimately "irrational."

Neuroeconomic studies have only deepened the rabbit hole with the recent discovery that distinct regions of the brain—a region associated with emotion and one associated with reason—are activated in response to potential for immediate reward versus potential for future reward, respectively.

We analyze the foundations of hyperbolic discounting through the lens of *rule-rationality*, an evolutionary approach to rationality that proposes that people do not maximize utility in each of their acts; rather, they follow rules of behavior that usually (e.g., on average), but not necessarily always, maximize utility. Under rule-rationality, we maximize over rules rather than acts (Aumann, 2008).

Classical decision theory presumes that, given constant information, an individual's choice of one option over another will not fluctuate as a function of *when* that choice is made. If utility is definitively maximized when choice α is chosen over β in period t, *ceteris paribus,* it should be maximized when α is chosen over β in period $t + 1$. To be safe, even if we allow for individuals to "change their minds," it does not explain the *systematic* deviation from time-consistent choice that results from hyperbolic discounting. Rule-rationality, however, involves the application of rules to decisions—rules that evolve with the singular purpose of maximizing *fitness*, one's expected number of offspring. Therefore, it is possible that the alleged inconsistent choices that result from hyperbolic discounting are not inconsistent at all, but rather just governed by different rules.

We will show that evolution has instilled a rule that favors immediate reward over larger but later reward. However, the same variable that has enforced such a rule has simultaneously dictated against its extension to decisions not involving immediate reward: *unpredictability*. As will be discussed further below, individuals execute rules by means of a *mechanism*, which expresses a rule only indirectly. Indeed, the evolutionary advantage in foregoing choices that rely on relatively static environmental conditions provides grounds not only for the development of an impulsive mechanism, but also for not extending such a mechanism to decisions not involving immediate reward. I analyze hyperbolic discounting within Robert Aumann's rule-rationality framework to show that it has evolved precisely *because* it is rational.

II. Hyperbolic Discounting—History and Overview

There is a trend in the economic literature to replace the constant discount utility function with one characterized by hyperbolic discounting. Classically, economists assumed a discounting of future reward by a factor of $(1 + k)^{-t}$, where *k* is the constant discount rate per unit of time and *t* is the duration of delay. This is known as "exponential discounting." The prevailing view today, however, is that exponential discounting does not adequately account for the data. It assumes that individuals discount future reward by a fixed percentage for each period of time delay. For example, an individual who discounts at 10% per year would be indifferent between $100 today and $110 in one year from now; he would similarly be indifferent between $100 in one year from now and $110 in two years from now. Future rewards are discounted as a function only of the duration of delay and constant rate of discounting.

But studies, as well as our own thought experiments, show that people do not in fact discount exponentially. They seem to discount at higher rates when the options are near in time and lower rates when they are more distant in time. $100 today is usually preferred to $110 in a month from now; but $100 in a year is rarely preferred over $110 in a year and a month.

Thus, many economists have adopted hyperbolic discounting as a replacement. Hyperbolic discounting decreases future reward by a factor of $(1 + kt)^{-\beta/\alpha}$, where α and β are greater than zero. The formula is that of a generalized hyperbola and was originally adopted by psychologists to model such preferences (see Lowenstein and Prelec, 1992). Simpler versions of hyperbolic discounting have since been proposed and have gained wide popularity as well. Researchers have proposed a "quasi-hyperbolic discount function"[1] involving two parameters, β and δ, in which reward obtained in periods 0, 1, 2, 3, … are discounted by 1, $\beta\delta$, $\beta\delta^2$, $\beta\delta^3$…, respectively (see Phelps and Pollak, 1968; Laibson, 1997).

[1] $u(x_0, x_1, \ldots, x_t, \ldots) = v(x_0) + \beta\Sigma_{t = 1, 2, \ldots}\, \delta^t v(x_t)$.

Note, however, that the "abandonment" of constant discounting for hyperbolic discount functions has not gone without criticism. Rubinstein (2003) reminds us that there are an infinite number of functional forms consistent with the findings that support hyperbolic discounting. Rubinstein concludes that the data—including data resulting from his own experimentation—eludes the hyperbolic discounting utility function as it does other functional forms (Rubinstein, 2003, 2006).

Suffice it to say, hyperbolic discounting has gained large support among economists, and in fact, it *does* seem to better fit the empirical data. Studies comparing the hyperbolic discount structure to the exponential framework have overwhelmingly found in favor of the former.[2] Indeed, the vast evidence favoring the hyperbolic structure has prompted its wide acceptance—and with it, revival of the discussion of "time-inconsistency" (or "dynamic-inconsistency") dating back to Strotz (1956).

Consumers seem to be "of two minds" when it comes to intertemporal consumption decisions (Laibson, 1998). An individual makes the decision, for example, to diet rather than to enjoy eating cake; but as the moment of sacrifice (cake avoidance) approaches, one tends to reverse his preference and choose to eat the cake rather than diet. After the cake has been eaten, however, the individual often returns to his original preference of dieting.

Thus, the discovery that people discount hyperbolically is not one that simply calls for replacement of the constant discounting utility function; it requires an overhaul that addresses an individual player's numerous "minds" with respect to intertemporal decisions. The most prevalent of such models entails the image of an individual as a "composite of autonomous temporal selves"—selves who interact as players in a finite-horizon dynamic game (see Laibson, 1998).

Indeed, hyperbolic discounting has been applied to explain a wide variety of behavioral phenomena involving self control—phenomena ranging from addiction to procrastination, from undersaving to undercommitment.

The image of a single individual as a composite of autonomous temporal selves is enhanced by recent neuroeconomic studies that reveal an *actual* duality of minds with respect to intertemporal decisions.

Researchers used functional magnetic resonance imaging (fmri), a procedure that images brain activity, to find that areas of the brain associated with emotion dominate decisions involving the possibility of immediate reward, while areas of the brain associated with logic dominate

---

[2] See Laibson (1998), citing competitions testing the use of each form. See also Chung and Herrnstein (1961); Ainslie (1992); Kirby (1997).

decisions involving the same reward set, but delayed. When subjects were presented with the option of immediate reward but chose delayed reward, the brain's logic centers dominated the decision, while a choice of immediate reward displayed no such domination, and instead, a slight trend toward heightened emotional activity (see McClure et al., 2004).

Researchers have concluded that as proximity to reward narrows, impulsive choices resulting from areas of the brain associated with emotion "wins out" over the abstract-reasoning parts of the brain. It is reported that preference reversals that occur as a result of hyperbolic discounting are, "[i]n classic economic theory … irrational because people are inconsistent in their treatment of the [constant] time delay." One of the researchers, Samuel McClure, remarked that "[t]here are two different brain systems, and one of them kicks in as you get really proximate to the reward."

David Laibson, who also took part in the study, explains as follows:

> Our emotional brain has a hard time imagining the future, even though our logical brain clearly sees the future consequences of our current actions …. Our emotional brain wants to max out the credit card, order dessert and smoke a cigarette. Our logical brain knows we should save for retirement, go for a job and quit smoking. To understand why we feel internally conflicted, it will help to know how myopic and forward-looking brain systems value rewards and how these systems talk to one another.[3]

Of course, it is true that neither inconsistency nor "conflicting selves" is a concept that sits well with the classical assumption of rational choice. It is therefore tempting to simply categorize hyperbolic discounting as an irrational behavior, as numerous economists have. But is there not a deeper explanation, one that might reconcile concepts of rationality and preference reversal?

We answer in the affirmative. We apply Aumann's *rule-rationality* to present a straightforward explanation of the rational foundations of hyperbolic discounting and its characteristic preference reversal.

III. Rule-Rationality

Robert Aumann's *Rule-Rationality Versus Act-Rationality* (2008) proposes that people do not maximize utility in each of their acts; rather, they follow rules that optimize in the aggregate. We maximize utility over rules rather than acts.

[3] Princeton University (October 14, 2004), Study: Brain Battles Itself over Short-Term Rewards, Long-Term Goals.

Importantly, individuals only rarely adopt rules consciously. They often do so as the result of genetic or memetic evolutionary forces, or as the product of a learning process. Further, individuals execute rules by means of a *mechanism*, which expresses rules only indirectly (p. 2).

Let us take the Ultimatum Game as an example. Two players, the *proposer* and the *responder*, must divide 100 Euro between themselves. The *proposer* begins by proposing a division. The *responder* must then accept the proposed division or reject it. If he accepts the division, each player receives the proposed share; if he rejects the division, each player receives nothing, and the game ends, not to be played again.

Experimentation with the Ultimatum Game repeatedly results in departures from rational choice predictions. In one such experiment, two individuals (students) were seated at computer stations in separate rooms. The individuals were not able to communicate and, after the game is played, were to leave through separate doors without seeing each other or learning of each other's identity. The *proposer* was to provide his offer by simply entering a numerical value into the computer; the *responder* was to reply simply with a "yes" or "no."

One might expect the *proposer* to have offered a division such that the *responder* receives some non-negligible amount—say, 90 Euro for the *proposer* and 10 Euro for the *responder*—and for the *responder* to accept the offer. Although the division may seem "unfair" to the *responder*, so long as the offer provides him with a non-negligible amount, why would he not accept? In the example above, the *responder's* option is to receive 10 Euro or nothing. He knows that the game will end and will not be repeated; therefore, he has little incentive to "teach a lesson" to his opponent. In fact, he will not even see, communicate with, or know the identity of his opponent.

But the findings do not confirm such expectations. Repeatedly, the *proposer* offered approximately a 65-35 split (in favor of himself), and when the proposal offered the *responder* substantially less—say an 80-20 split—the *responder* often rejected the proposal. The *responder* chose to receive nothing rather than 20 Euro!

Numerous explanations have been offered, including those involving feelings of pride, revenge, etc.—all legitimate sources of utility. One may, for example, gain positive utility from preserving his pride or exacting "revenge" upon the *proposer* by rejecting his offer. As Aumann (2008) explains, such positions are perfectly logical; however rule-rationality offers a deeper explanation—an accounting of such emotions in terms of more fundamental needs.

Aumann suggests that while the decision of the *responder* to reject an 80-20 offer may not be *act-rational*, it is indeed *rule-rational*. Individuals, according to rule-rationality tend to adopt certain rules that govern real-world proposals. Such a rule might demand the rejection of unfair offers. Although the aforementioned Ultimatum Game experiment did not allow for reputational effects, real-world interaction usually *does*—and a reputation of accepting only "fair" offers will likely prompt a more even-handed approach by real-world *proposers*. Application of such a rule may result in occasional deviations from act-rationality; but the rules evolve precisely because its application is *usually* act-rational.

Further, an individual may be completely unaware that he is executing a rule. According to rule-rationality, rules are executed via a *mechanism*. Aumann explains that in the Ultimatum Game example above, the mechanism might consist of a mix of emotions such as pride and a desire for revenge. These emotions are likely to have evolved as a response to unfair offers particularly *because* they execute a fitness-maximizing rule.

Fundamental to the rule-rationality paradigm is the relationship between evolution and rationality. As noted in *Rule-Rationality Versus Act-Rationality*, the connection between evolution and rationality has been recognized for decades.[4] Aumann describes rationality as a *product* of evolution: he asserts that evolution is the "fundamental driving force" behind ordinary utility-maximizing rationality. "Rationality has evolved, alongside of physical features like eyes, stomachs, limbs, and breasts, *because it maximizes fitness*" (p. 4).

The connection essential to extending the relationship between evolution and rationality beyond analogy is that between utility and fitness. Rationality maximizes utility while evolution maximizes fitness. But, as Aumann points out, man does not shape his utility—an expression of preference, by the usual definition—with the sole consideration of increasing his fitness, the number of his offspring. "[A]n obese person craving another piece of chocolate maximizes his 'utility' by eating it, but surely not his fitness" (p. 16).

Aumann argues that defining preference in terms of how one would choose if faced with an option, and defining rationality (i.e. act-rationality) as acting so as to maximize utility—which is defined by preference—implies that one's actual choice must be rational, as it is the preferred one. But this implies that *all* behavior is rational, and that rationality is, in a sense, a tautology.

Aumann concludes that to make sense of the literature, one must define utility in terms of basic desiderata such as time, money, family welfare, life, health, and food—items closely related to fitness. With such a definition, "an act that maximizes utility is then indeed act-rational, and a rule that usually maximizes it, rule-rational" (p. 17).

Like the behavior witnessed in Ultimatum Game experiments, hyperbolic discounting may be explained in a number of ways. Rule-rationality provides a framework through which we may examine the rational basis of hyperbolic discounting in fundamental terms, and in terms of its evolutionary foundations.

[4] See p. 3 (citing Maynard Smith and Price, 1973, Dawkins 1976).

IV. The Rational Foundations of Hyperbolic Discounting

Hyperbolic discounting is a unique behavioral pattern in that it sits at the intersection between reasoned decision making and impulse. It therefore begs for a theory that can speak to both aspects. Classical definitions of rational choice do not adequately account for the impulsive behavior that characterizes a hyperbolic discounter who is near in time to potential reward. Behavioral economists, on the other hand, have shown little hesitation in simply defining such behaviors as irrationalities.

Challenges to the rationality assumption, for example, have often confused economic "rationality" with "reasoning." The two, however, are far from synonymous. Rational behavior means acting in a manner that is consistent with one's utility function. The evolutionary context in which rule-rationality is embedded underscores the point that reason, in the sense of logical deduction, constitutes only a single aspect of rational choice.

Evolution has imposed distinct forces on our decision-making process, and a range of mechanisms to implement them. We claim that within the evolutionary context set forth above, reason and impulse are both entirely valid forces of rational choice. Both are products of evolution in the sense of fitness maximization, and both valid with respect to their consistency with utility maximization.

Hyperbolic discounting encompasses a preference reversal that is commonly attributed to our impulsive tendencies. It is framed as a battle between impulse and reason, and one in which the former emerges victorious in decisions involving immediate reward.

But is this phenomenon best characterized as a battle, or has evolution imposed a mechanism by which a fitness-maximizing rule is executed?

To start, consider an example addressed in Aumann's *Rule-Rationality Versus Act-Rationality* in which the effect of evolution is explicit: the gene for altruism. For many years, displays of altruism have puzzled economists. Why do people behave altruistically in single-shot games? It is one thing to act with altruism when one believes that the behavior will be returned in turn, even if no explicit agreement exists. It is another to do so when it is clear that "repayment" is out of the question. Similarly, behavioral economists have spent years researching the "Dictator Game," an experiment in which one player decides the distribution of a fixed sum of money between himself and another. It is clearly explained to the subject that the game will be played only once and that afterward, the player will leave with the distributed money without further consequence, and without ever meeting the player's counterpart. Interestingly, many players who decide on the distribution (the "dictators") allocate a percentage of the money to their counterparts (the "recipient" players). Why should this occur? What is gained by acting altruistically in such a circumstance?

Indeed, the Dictator Game has been cited by behavioral economists as evidence against the rationality assumption. Others categorize behaviors such as altruism and revenge simply as "other-regarding preferences." Many economists, including Aumann, do consider altruism a perfectly legitimate source of utility. As mentioned above in the current paper, there is little reason to relegate "psychological utility" to a lesser validity. Aumann (2008) addresses the issue of altruism as a prime example of rule-rationality. In repeated interactions, altruism is often act-rational, but in one-time encounters, altruism may be rule-rational, even when it is not act-rational. This is so because "[r]ather than keeping accounts of who helped whom when, it may be simpler just to be generous as a rule." As such, altruism can be seen as a mechanism for achieving cooperation, a strategy that is often—perhaps usually—a good idea (pp. 11-12).

Now here's the twist: researchers have recently found a gene that is associated with altruism in the Dictator Game—a concrete physiological basis for altruism. Perhaps this gene evolved in humans precisely because cooperation is a rational behavior, repeated interaction or not. As Aumann put it, "It is all well and good to speak airily about 'other-regarding preferences,' but when you have a gene staring you in the face, you've got to ask yourself, from where did this come? How—and why—did this gene *evolve*?"

Although not as explicit as the gene for altruism, consider the neuroeconomic research cited above, in which researchers demonstrated a discrepancy in brain activity when immediate reward is involved in a given decision versus a choice involving rewards more distant in time. Let us encounter the issue head-on: why respond with impulse when reward-potential is close in time, but choose with reason when it is not?

Rule-rationality suggests that individuals may display altruism in a one-shot game not because individuals are inherently irrational, but rather because evolution has *selected* for the survival and flourishing of this behavioral characteristic—and it has done so precisely *because* it is optimal in the aggregate. It is a mechanism for achieving fitness-maximizing behavior: cooperation!

As with altruism in the Dictator Game, it is tempting to immediately throw one's arms in the air and yell "irrationality" at the discovery of hyperbolic discounting. Indeed, many economists have fallen to such temptation. But other explanations do exist, and rule-rationality supplies a lens that brings the confusion surrounding hyperbolic discounting into focus.

It is not difficult to recognize the evolutionary advantage for valuing immediacy of reward—from favoring an immediate payoff even at the cost of sacrificing a larger, albeit later, payoff. Indeed recent studies confirm our intuition. Scientists have shown that "[impulsive behavior] may have evolved because in the wild, snatching up small rewards like food morsels rather than waiting for something bigger and better to come along can lead to getting more rewards in the long run."[5] Researchers have found that impulsive strategies (i.e. immediate-reward centric strategies), implemented by observed birds, yielded as much or more food in the long run than strategies involving holding out for larger rewards.

Researchers were unable to train birds to forego immediate reward in order to obtain a larger, later reward. The results have been explained as the birds' instinctive realization that a reward delayed may be a reward foregone, since conditions can change as the bird is waiting. "Animals, I think, come with a hardwired rule that says, 'Don't look too far in the future,'" explained David Stephens, a professor of ecology, evolution, and behavior at the University of Minnesota. "Being impulsive works really well because after grabbing food, they can forget it and go back to their original foraging behavior. That behavior can achieve high long-term gains even if it's impulsive."

The findings, according to Stephens, may apply to humans, because taking immediate reward may have provided an important advantage to our foraging ancestors, as it does for other foragers.

Indeed, the application of such advantages remains intuitive today, and the neuroeconomic results cited above suggest that our bodies have accounted for such advantages by developing neural mechanisms to address choices involving the possibility for immediate reward. Circumstances change rapidly, and so our wiring has evolved to override, in a sense, certain neural systems associated with abstract reasoning when faced with an opportunity to receive immediate reward. Or, it is possible that we have evolved in the first instance to act impulsively, and only more recently have begun to incorporate abstract reasoning into the equation.

Either way, it seems that humans may well have developed a mechanism, displayed in impulsive behavior, for executing a rule, *a la* rule-rationality, that favors immediacy of reward. In terms of *evolutionary fitness*, foregoing a larger, later payoff in favor of a smaller, immediate payoff may not always be act-rational—as we well know, impulsive behavior has a strong destructive potential. Nevertheless, evolutionary forces may have induced such a mechanism precisely because snatching up immediate reward, even at the cost of sacrificing larger, future reward, is *rule-rational*—it is a strategy that, in the aggregate, maximizes fitness!

This explanation is evidenced by behavioral studies such as Stephens's, and even neuroeconomic studies such as those described above.

---

[5] University of Minnesota (December 6, 2004), Impulsive Behavior May Be Relict of Hunter-Gatherer Past.

Of course, neither the aforementioned behavioral study nor rule-rationality implies necessarily that immediate reward should be preferred over a larger but delayed reward. Indeed, Stephens cites a study finding that children who are good at waiting for reward "generally do better in life." This is not the issue we are addressing—we do not, by any means, claim that one should, as a rule, impulsively jump at immediate reward rather than investing in larger future reward. Rather, we intend to explain the evolutionary foundations of *valuing* immediacy of reward—and valuing immediacy of reward to the extent that discounting functions take hyperbolic form.

Indeed, people do not place an infinite value on immediacy. If rather than offering a choice of $10 today versus $11 tomorrow, I offered $10 today versus $20 tomorrow, a subject is likely to be more inclined to forego the immediate reward; if I offer $50 tomorrow, even more so.

To be clear, we are explaining a behavioral pattern purported by many to be irrational as, rather, an evolutionary mechanism that executes advantageous behavior. Now we are in a position to better understand the preference reversal that characterizes hyperbolic discounting.

This phenomenon can be understood through rule-rationality. Specifically, the evolutionary mechanism relating to impulsive action applies only to choices involving immediate reward; it does not extend to future decisions.

Evolutionary forces have instilled a rule that favors immediate reward over larger, but later reward. It follows from the above discussion, however, that the rule would take effect only when the potential reward is immediate. The entire purpose for the rule would be reversed if the mechanism were applied to choices far off in the future.

Specifically, the rule applies only to choices involving immediate reward likely for the same reasons that such a rule evolved in the first instance: to account for rapidly changing environmental conditions. The evolutionary advantage in foregoing choices that rely on relatively static environmental conditions provides grounds not only for the development of an impulsive mechanism, but also for not extending such a mechanism to decisions not involving immediate reward.

V. Games Researchers Play

In light of the experimental results that have emerged from the field of behavioral economics, one might be inclined to question the current relevance of the above analysis with respect to the issue of rationality. Do advantages gained by our foraging ancestors in valuing immediacy of reward apply in today's modern world? Does the evolution of a gene for altruism really provide an advantage in today's society? Perhaps it can be argued that rules that provided an evolutionary advantage to our ancestors actually impose a disadvantage in modern civilization—perhaps the vestiges of a state of rationality from a past time have created the seeds of today's irrational behavior. The experimental data discussed above certainly suggest as such—but justifiably so?

We argue that many of the recent psychological experiments are, in a sense, *rigged* to result in data that implies irrationality. Researchers often play "tricks" on their subjects in the sense that they present a fantasy circumstance for which the subject is trained by real-world circumstances to act in a particular fashion—a fashion that is optimal generally, but not as applied to the given fantasy conditions.

In terms of rule-rationality, years of real-world circumstances have caused the evolution of certain mechanisms the purpose of which is to execute rules of behavior that are optimal in the aggregate. Researchers have exploited these rules, for the purpose of "proving" human irrationality, by selecting unusual cases—cases that rarely reflect real-world circumstances—for which the rule fails to result in the optimal outcome. Let us take two examples.

As discussed above, researchers have found that subjects playing the Ultimatum Game rejected as much as twenty Euro for reasons perhaps involving feelings of pride, revenge, etc., even though they were well aware that the game would not be repeated and they would not meet their counterparts. Researchers have published countless articles on such "anomalous" results. But should these results be so unexpected?

In the "real world," circumstances *do* repeat among individuals; players almost always *do* meet or know the individuals with whom they interact. Researchers have devised a payoff scheme similar to some real-world conditions, but have altered crucial conditions, such as repeated interaction. But individuals are *trained* to respond to the real-world: they execute rules—often subconsciously—to optimize their payoff in response to offers they receive under real-world conditions.

So is it surprising that subjects in the Ultimatum Game have rejected offers they are likely to view as unfair in real-world scenarios? In the context of rule-rationality, it is not surprising at all. In fact, it is rule-rational to do so.

Similarly, subjects of the Dictator Game experiments can be seen offering a percentage of the total sum they receive to their counterparts, although they seemingly gain nothing from so doing. Here, the exploitation is even more overt.

Real-world circumstances have demanded cooperation to the extent that individuals have developed a gene associated with altruism. The gene emphasizes the long process of "training," or evolutionary development that stands behind the real-world circumstances to which such mechanisms apply.

Thus, when researchers instruct subjects to divide a sum of money as in the Dictator Game, it should be expected that some of the subjects will offer a percentage to their respective counterparts. Is it irrational to do so? As discussed at length above, even if the behavior is defined as *act-irrational*, the gene for altruism has evolved, pursuant to rule-rationality, precisely because it provides a mechanism that delivers rule-rational behavior. In the real world, altruism *does* pay off—it results in cooperation!

Finally, let us consider the hyperbolic discounting experiments in which, for example, researchers offer subjects the choice of $10 today or $11 tomorrow, and a parallel choice of $10 in a month or $11 in a month and a day. As discussed above, the resulting data shows that most subjects choose the smaller but immediate payoff in the former option, but the larger later payoff in the latter. The results have been labeled dynamically inconsistent and have further fueled the claim that individuals cannot be considered rational, as conventionally assumed.

Throughout man's history, however, we have been taught that $10 today versus $11 tomorrow *is not* equivalent to $10 in a month versus $11 in a month and a day: one choice involves immediate reward and the other does not—and our neural makeup has *accounted* for such distinctions.

Researchers have devised experiments that trigger one's instinct with respect to his non-reliance on a static environment. Our usage of the word "instinct" is appropriate because the experiments do not account for the real-world condition that foregoing $10 today for the sake of gaining $11 tomorrow may result in a payoff of $0 rather than $11. Indeed, the subject himself often does not presume that should he pick $11 tomorrow over $10 today, there is a possibility that he will receive no payoff at all,[6] just as he likely does not presume—at least consciously—that the choice of $10 in a month versus $11 in a month and a day might fail to remain static. It is his *instinct* that is at issue: his subconscious ability to discern a choice involving immediate reward, to associate such a choice with its inherent risks, and to trigger a mechanism for impulsive behavior.

As in the Ultimatum and Dictator Game experiments, researchers have exploited their subjects' historical training by replacing real-world conditions with fantasy games in which elements such as risk are ignored. But subjects cannot simply ignore subconscious mechanisms that have evolved to respond to real-world conditions—that is, to *rationally* respond to such conditions.

---

[6] Of course, exceptions to this do exist—for example, subjects of experiments on remote indigenous tribes, as well as skeptics even within major modern cities.

Games devised to "trick" the mind into choosing an act-irrational behavior, as with certain unusual real-world cases, may result in suboptimal choices, but choices which are nevertheless *rule-rational*. In this sense, many of the arguments and empirical results that serve as the foundation of the "irrationality movement" in the economic literature fail to hold water.

VI. Conclusion

Defining rationality in the context of evolution requires a dynamic perspective. We are, of course, not perfectly evolved creatures, and evolution itself continually spawns new meanings of rationality, new life-demands to be satisfied, and new mechanisms to relate behavior to utility. Times *do* change, and strategies that were optimal for our foraging ancestors may not always be optimal in today's modern world. Nevertheless, the evolutionary paradigm, and rule-rationality in particular, suggests that we take a perspective in which optimality, or rationality, is measured not by the payoff a behavior generates in any single application, but by its effect in the aggregate, over its application as well as over time.

Rule-rationality offers a framework for analyzing behavior in terms of its evolutionary constituents, and one that allows a deeper understanding of seemingly anomalous behavior all too often hastily branded as irrational.

We have analyzed hyperbolic discounting in the context of rule-rationality with the aspiration of uncovering its rational foundations. In so doing, we have interpreted the discussed neuroeconomic findings as evidence of a biological trigger that has evolved for the purpose of inducing a generally subconscious *mechanism*, displayed as impulsive behavior, that ultimately executes a rule and delivers rule-rational behavior. We are optimistic that advances in neuroeconomics, biology, and psychology will bring about further understanding of the biological and evolutionary foundations of such mechanisms.

Finally, we have emphasized that the rational foundations of hyperbolic discounting, as well as numerous other behaviors discussed, should not be relegated to a category of mere historical importance. Recent behavioral studies have shown us only that subjects do not necessarily optimize with every action in every circumstance—results that are consistent with rule-rationality. Many behavioral experiments have resulted in data that demonstrates a divergence from classical "rationality," but are characterized by fantasy conditions that simulate real-world circumstances in some respects while omitting certain critical aspects. These experiments often aim to "trick" subjects who have developed generally optimal responses to real-world conditions into applying such mechanisms to cases in which they are not optimal. In short, they exploit rule-rational behavior and its underlying mechanisms.

Hyperbolic discounting is one such behavioral pattern that has received much attention in the recent literature—and justifiably so, as it underlies an array of behavioral anomalies seen throughout today's society. Although hyperbolic discounting, and impulsive behavior in particular, may contain a certain destructive potential, it is likely that its evolutionary foundations are sound—and its application may well be as justified and rational today as it was for our foraging ancestors. Rule-rationality suggests as such, reminding us that we can learn a great deal by adopting a broader and more dynamic perspective of economic behavior than is prevalent in today's literature—a particularly potent reminder in the context of analyzing hyperbolic discounting, a behavior that emphasizes our dynamic treatment of choices over time.

References

Ainsle, G. (1992), *Picoeconomics*, Cambridge University Press, New York, NY.

Aumann, R. J. (2008), "Rule-Rationality Versus Act-Rationality," *Discussion Paper Series* dp497, Center for the Study of Rationality, The Hebrew University, Jerusalem.

Chung, S. and Herrnstein, R. (1961), "Relative and Absolute Strengths of Response as a Function of Frequency of Reinforcement," *Journal of the Experimental Analysis of Animal Behavior* 4, 267-272.

Kirby, K. N. (1997), "Bidding on the Future: Evidence against Normative Discounting of Delayed Rewards," *Journal of Experimental Psychology: General* 126, 54-70.

Laibson, D. (1997), "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics* 112(2), 443-477.

Laibson, D. (1998), "Life-cycle Consumption and Hyperbolic Discount Functions," *European Economic Review* 42(3), 861-871.

Lowenstein, G. and Prelec, D. (1992), "Anomalies in Intertemporal Choice: Evidence and Interpretation," *The Quarterly Journal of Economics* 107, 573-597.

McClure, S. M., Laibson, D. I., Lowenstein, G., and Cohen, J. D. (2004), "Separate Neural Systems Value Immediate and Delayed Monetary Rewards," *Science*, 306(5695), 503-507.

O'Donoghue, T. and Rabin, M. (1999), "Doing It Now or Later," *American Economic Review* 89(1), 103-124.

Phelps, E. S. and Pollack, R. (1968), "On Second Best National Savings and Game-Equilibrium Growth," *Review of Economic Studies* 35(2), 185-199.

Princeton University (October 14, 2004), "Study: Brain Battles Itself over Short-Term Rewards, Long-Term Goals," *available at* http://www.princeton.edu/pr/news/04/q4/1014-brain.htm.

Rubinstein, A. (2003), "'Economics and Psychology'? The Case of Hyperbolic Discounting," *International Economic Review* 44(4), 1207-1216.

Rubinstein, A. (2006), "Dilemmas of an Economic Theorist," *Econometrica* 74(4), 865-883.

Strotz, R. H. (1956), "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies* 23, 165-180.

University of Minnesota (December 6, 2004), "Impulsive Behavior May Be Relict of Hunter-Gatherer Past," *available at* http://www1.umn.edu/news/news-releases/2004/UR_RELEASE_MIG_1997.html.