Rational Choice Theory

Necessary but Not Sufficient

R. J. Herrnstein *Harvard University*

ABSTRACT: A case is presented for supplementing the standard theory of rational choice, according to which subjects maximize reinforcement, with a theory arising from experiments on animal and human behavior. Data from these experiments suggest that behavioral allocation comes into equilibrium when it equalizes the average reinforcement rates earned by all active response alternatives in the subject's choice set. This principle, called the matching law, deviates from reinforcement maximization in some, but not all, environments. Many observed deviations from reinforcement maximization are reasonably well explained by conformity to the matching law. The theory of rational choice fails as a description of actual behavior, but it remains unequaled as a normative theory. It tells us how we should behave in order to maximize reinforcement, not how we do behave.

We start with a paradox, which is that the economic theory of rational choice (also called optimal choice theory) accounts only poorly for actual behavior, yet it comes close to serving as the fundamental principle of the behavioral sciences. No other well articulated theory of behavior commands so large a following in so wide a range of disciplines. I will try to explain the paradox and to present an alternative theory. The theory of rational choice, I conclude, is normatively useful but is fundamentally deficient as an account of behavior.

Rational choice theory holds that the choices a person (or other animal) makes tend to maximize total utility, where utility is synonymous with the modern concept of reinforcement in behavioral psychology. Because utility (or reinforcement) cannot be directly observed, it must be inferred from behavior, namely, from those choices themselves. Rational choice theory is thus a rule for inferring utility: It says that what organisms are doing when they behave is maximizing utility, subject to certain constraints. Rational choice theory is also used normatively, as a way of assessing whether behavior is, in fact, optimally gaining specified ends, and if not, how it should be changed to do so. The distinction between descriptive and normative versions of rational choice theory is fundamental to the theme of this essay.

The theory of rational choice seems to stand in relation to the behavioral sciences as the Newtonian theory of matter in motion stands to the physical sciences. It is held, by its proponents, to be the law that behavior would obey if it were not for various disruptive influences, the behavioral analogues of friction, wind, measurement error, and the like.

Not just economics, but all the disciplines dealing with behavior, from political philosophy to behavioral biology, rely increasingly on the idea that humans and other organisms tend to maximize utility, as formalized in modern economic theory. In accounts of governmental decision making, foraging by animals, the behavior of individual or collective economic agents, of social institutions like the criminal justice system or the family, or of rats or pigeons in the behavior laboratory, it has been argued forcefully that the data fit the theory of rational choice, except for certain limitations and errors to which flesh is heir. The scattered dissenters to the theory are often viewed as just that—scattered and mere dissenters to an orthodoxy almost as entrenched as a religious dogma.

How can anyone plausibly subscribe to the descriptive theory of rational choice in the face of the reality that organisms often behave against self-interest? Even some rational choice theorists procrastinate and suffer from other human frailties. They may overeat, smoke, drink too much, and make unwise investments, just like the rest of us. People may behave altruistically at some personal sacrifice. Martyrs are just rare, not unknown. Neither the existence of unwise nor altruistic actions evidently wounds the descriptive theory of rational choice for its most committed adherents.

A resistance to ostensibly contrary data is not unique to rational choice theory. It has often been observed that scientific theories evolve to cushion themselves from the hard knocks of data; neither rational choice theory nor the alternative theory to be proposed here is an exception to this generalization.

But that general resistance to counterevidence is not the only reason rational choice theory endures. Behavior that might seem irrational because it is not guided by obvious self-interest is sometimes explained in rational choice theory by invoking whatever source of utility is needed to rationalize the observed behavior. This is pos-

 0003-066X/90/$00.75
Vol. 45, No. 3, 356–367

sible within the theory because utility, which is subjective, differs from objective value. There is, in principle, no constraint on utility other than that imposed by the behavior from which it is inferred. In principle, nothing prevents inferring utilities that lead to self-damaging or altruistic behavior, for example. A similar strategem is available to reinforcement theorists, who are also free to infer reinforcement from the observed behavior.

We may, for example, be optimizing subjective utility (or reinforcement) by eating ice cream and red meat and smoking dope, even though we are, and know we are, harming ourselves. Some people give up a great deal, objectively speaking, for the subjective utilities they are presumably deriving from cocaine or alcohol, including shortening their lives and decreasing the quality of their lives. The things that organisms strive to obtain or to eliminate are taken as givens by the theory. When rational choice theorists say, "De gustibus non est disputandum," they mean it (Stigler & Becker, 1977). Rationality, in this modern version, concerns only revealed preference.

Not only are utilities subjective, says the theory of rational choice, but so are the probabilities by which they can be discounted by uncertainty. People often act as if they overestimate low, but nonzero, probability outcomes and underestimate high probability outcomes, short of certainty. They may worry too much about, and pay too much to insure themselves against, low-probability events such as airplane accidents. People insure their cars against improbable losses, then, with abandon, run red lights on heavily traveled city streets. After working hard to earn their pay, they buy lottery tickets with infinitesimal odds of winning. Instead of objective probabilities, it has been proposed that utility theory must take into account subjective weights, bearing complex, as yet unexplained, relations to objective frequencies.

The subjectivity of utility is motivational. The subjectivity of probability is cognitive. Rational choice theorists invoke other psychological complications beyond these, having to do with limitations in organisms' time horizons, knowledge, capacities for understanding complexity, and so on. Acknowledging those limitations, while saving the theory, is like the postulation of epicycles in planetary astronomy, in either case smoothing the bumpy road between facts and theory. The question is whether the epicycles of rational choice theory are protecting a theory that inhibits understanding or advances it, whether the correct analogy is Ptolemy's geocentric theory or Copernicus's heliocentric theory, each with its own epicycles.

As a descriptive theory, rational choice theory survives the counterevidence by placing essentially no limit of implausibility or inconsistency on its inferred utilities and also by appealing to the undeniable fact that organisms may calculate incorrectly, be ignorant, forget, have limited time horizons, and so on. Other lapses of rationality, as they are illuminated by the numerous ingenious paradoxes of choice research, are often swiftly absorbed by the doctrine of rational choice, at least in the eyes of its most devoted followers. Those odd, obscure, or shifting motives and those errors of calculation and time perspective aside, we are all rational calculators, the theory says.

Rational choice theory also survives because it has several genuine strengths, beyond its indisputable value in normative applications. First, rationality accords with common sense in certain simple settings. For example, consider a choice between \$5 and \$10, no strings attached. Any theory of behavior must come up with the right answer here, where there seems to be no issue of obscure motives, or of errors of reckoning, remembering, knowing, and so on. Assuming only that more money has more utility than less money, rational choice theory does come up with it. To argue against rationality as a fundamental behavioral principle seems to be arguing against self-evident truth.

Second, rational choice theorists have formalized utility maximization, reducing it to its axiomatic foundations. Many of the most brilliant theoreticians are drawn to this part of the behavioral and social sciences, for here is where their powerful intellects shine most brightly, addressing questions of formal structure, not distracted by the fuzziness of motivation or the messiness of data. Some rational choice theorists admit that the theory is wrong, but they see no good reason to give up something so elegantly worked out in the absence of a better theory. Many rational choice theorists evidently believe that no theory could simultaneously describe behavior better than, and be as rigorous as, rational choice theory. Real behavior, they seem to believe, is too chaotic to be rigorously accounted for with any precision.

The foundations of rational choice theory have, however, lately been under attack. Experimental findings by many decision researchers (e.g., Kahneman, Slovic, & Tversky, 1982) have undermined the descriptive form of the theory by discovering choice phenomena that are consistent with (or at least not inconsistent with) principles of cognitive psychology, but inconsistent with rationality as commonly construed. Bombarded by these data, the unifying concept of rational choice may give way to a set of psychological principles, none of which is of comparable breadth, but which, in the aggregate, will account for actual behavior better than the global assumption of rationality (an approach exemplified in a recent textbook by Dawes, 1988).

Theoretical challenges also abound. It has been re-

I gratefully acknowledge the large contributions to this work of William Vaughan, Jr., Drazen Prelec, Peter de Villiers, Gene Heyman, George Ainslie, James Mazur, Howard Rachlin, and William Baum, all colleagues now or earlier. In other publications, interested persons may find more precision and detailed accounts of the data (e.g., Herrnstein, 1970, 1988; Vaughan & Herrnstein, 1987; and, especially, Williams, 1988). Several anonymous reviewers and Associate Editor Donald Foss deserve thanks for uncommonly helpful comments on an earlier version of the article. I owe thanks, too, to the Russell Sage Foundation for support and an environment during the academic year 1988–1989 that provided an opportunity for a study of the relations between economic and psychological theories of individual behavior.

Correspondence concerning this article should be addressed to Richard J. Herrnstein, Harvard University, William James Hall, 33 Kirkland St., Cambridge, MA 02138.

peatedly suggested that it is not individual behavior that satisfies principles of rationality, but natural selection (e.g., Frank, 1988; Hirshleifer, 1982; Margolis, 1987). Evolution, guided by natural selection, endows individuals with behavioral rules of thumb that may be individually suboptimal, but that in the aggregate, approximate optimality in some sense (Heiner, 1983; Houston & McNamara, 1988). A few theoreticians (e.g., Luce, 1988, 1989; Machina, 1987), drawing mainly on the paradoxes of choice in the face of uncertainty (e.g., the familiar Ellsberg and Allais paradoxes, discussed in Dawes, 1988), have been exploring the possibility of relaxing one or another of the axioms of rationality while retaining the rest of the formal theory.

At least a few (and perhaps many) economists and other social scientists would, at this point, defend rational choice theory only in its normative form and would agree that the descriptive form has lost its credibility in the face of too many "anomalies" of individual behavior—too many epicycles, in other words. For many of these theorists, there is a theoretical vacuum as yet unfilled. One can predict a surge of new theories to fill the void. In this article, I will attempt to fill a part, if not all, of the vacuum with a theory arising out of the experimental analysis of behavior.

The advantages of the present theoretical alternative are that it accords no less well with common sense than rational choice theory, that it lends itself to as rigorous a formal structure, that it has extensive empirical support, and that it is consistent with many of the irrational behaviors we actually observe in ourselves and others. The primary disadvantage, which may or may not prove to be decisive, is that the large experimental literature on which it is based comes mainly, though not exclusively, from studies of animal rather than human subjects.

Some Systematic Irrationalities

The weaknesses in rational choice theory are uncovered by systematic inconsistencies in behavior, which can sometimes be graphically illustrated by asking people to solve riddles. Their solutions may betray the inconsistencies. I will consider two riddles and one experiment that point toward the alternative theory to be developed here. However, even in advance of an account of the theory I am proposing, the riddles and the experiment show that something goes wrong when people are asked to make certain kinds of choices.

Suppose a person is asked to imagine winning a lottery and is given a choice between $100 tomorrow and $115 a week from tomorrow.[1] Whichever the person chooses (only hypothetically, because no money is given), the money is said to be kept in escrow by a Federal Reserve bank, then delivered by bonded courier. Now the person is asked to choose one. When I present a problem like this, a fair proportion of people choose the earlier but smaller payoff.

Now, those who choose the smaller payoff are asked to imagine winning another lottery and are given a choice between $100 tomorrow and $140,000 a year from tomorrow. Again, the Federal Reserve holds the money and delivers it on the schedule chosen. Everyone, I find, picks the more deferred but larger prize.

Finally, consider winning yet another lottery. The person is asked to choose between $100, 52 weeks from today or $115, 53 weeks from today. The Federal Reserve will do its usual fine job of holding and delivering the money. Most of the people who chose $100 in the first lottery switch to $115 here.

This natural pattern of choices violates the consistency implicit in rationality, and it does not seem to be a matter of obscure motives or of incidentally faulty arithmetic. Some more fundamental flaw in our decision making appears to be responsible. In the first lottery, those who choose $100 have, by their choice, revealed a discount rate of more than 15% per week. They have, in effect, said that they would be willing to forgo $15 (possibly even more) to get $100 a week sooner. If their discount rate was less than 15%, they would have chosen the later $115 over the earlier $100.

In the second lottery, the choice of $140,000 reveals a discount rate smaller than 15% per week, because when $140,000 is discounted at 15% a week for 52 weeks, the result is $97.69, less than the $100 the person could get by choosing the earlier payoff. As odd as it may seem, someone who thinks $100 tomorrow looks better than $115 deferred for a week should also think it looks better than $140,000 deferred for a year, if rationality prevailed.

From past experience, I know that some people, confronted with this lack of consistency in their choices, staunchly defend their rationality. They say things like, "I chose the smaller amount in the first lottery because another $15 isn't worth my thinking and worrying about for an extra week. An extra $139,900, however, is another matter altogether, well worth waiting a year for." It is because of such excuses that we add the third lottery, because here, too, one would be thinking about collecting another $15 for an extra week, yet most people find it worthwhile to do so when the week is a year deferred.

Nothing in rational choice theory can explain this curious inconsistency, yet it seems to be an example of an almost ubiquitous tendency to be overinfluenced by imminent events. The tendency toward impulsive, temporally myopic, decision making causes considerable grief, as we all know. Let us be clear about how the example exemplifies irrationality. The mere discounting of deferred consequences need not be irrational. If one postpones payment for work done or goods delivered, one will have to pay more than if one pays immediately. The sellers may calculate rationally that they are forgoing interest they could be earning or pleasure that they could be harvesting while the buyers hang on to the payment and garner the fun or the interest. Even if they are not calculating

[1] A version of the riddle using $100 and $120 was described by Herrnstein and Mazur (1987). No formal experiment has been done with either that version or the present one, but from informal observations, it is clear that many people succumb to the inconsistency described here. The quantitative features of the inconsistency have not been explored under controlled experimental conditions.

human beings, but rats or pigeons in a behavioral experiment, deferred consequences are likewise downgraded. Perhaps natural selection has already factored in something functionally equivalent to the rational consideration of foregone benefits.

In either case, if the discounting is rational, the rate should be fixed per unit time, barring gratuitous assumptions. Fifteen percent a week is 15% a week, now or next year, in the theory of rational choice. In the example, however, we reveal that we downgrade not only value, but also the rate at which we downgrade value. The discount rate may be 15% for next week, but for a week a year from now, the discount rate itself has shrunk so much that it leaves $115 looking better than $100 even though they are separated by a week.

Many problems of choice spread over time have a similar shape. Imagine, for example, that we could always select meals for tomorrow, rather than for right now. Would we not all eat better than we do? We may find it possible to forgo tomorrow's chocolate cake or second helping of pasta or third martini, but not the one at hand. People who are trying to lose weight pay dearly to spend time in dieting resorts ("fat farms"), where what they get for their money is losing the option of not eating on their own. The examples reveal our tendency to be inconsistent because of impulsiveness.

A poignant example of temporal myopia is provided by the discovery of genetic markers for Huntington's disease, a progressive, fatal disease of the nervous system. The disease is typically asymptomatic until early adulthood or middle age. It is caused by a single, dominant gene, so that an offspring of one parent with the disease faces a 50–50 chance of having it himself or herself. It is now possible for people facing this risk to find out early in life, with high accuracy, whether or not they carry the gene.

By far, most of the people at risk have declined to take the test (Brody, 1988). This reluctance would make sense within a rationalistic framework if it were the case that the negative subjective change from a 50–50 chance to a virtual certainty of the disease were larger than the positive subjective change from a 50–50 chance to a virtual certainty of no disease. That, however, is the reverse of the evidence described in the newspaper article just cited.

People who know they face an even chance of this fatal disease have typically already factored much of the worst possible news into their lives, by choices made about marriage, parenthood, occupation, and so on. If their fears are confirmed, there is an increment of sorrow, a resignation to a fate already played out in their minds, but no huge change in subjective state. The newspaper account says that bad news triggers no visible increment in psychopathology or need for tranquilizers. In contrast, those who get good news experience enormous joy and relief. Over time, their lives probably readjust to normality. But even given this dramatic asymmetry favoring positive subjective change over negative, few people take the test.

The Huntington's example is faintly echoed in what happens when we stand in water up to our knees at the beach on a hot day, knowing that relief is only a few moments away if we plunge in.[2] But, instead, we are daunted by anticipation of those icy first few seconds. It can be so hard to overcome this barrier that we give up and turn back to the hot beach. Sometime between when we first left the blanket on the beach and when we hesitate knee deep, the promise of relief has been swamped by the avoidance of the rapid drop of temperature.

Note that these examples resemble the lotteries described earlier, in that an immediate consequence (e.g., the pleasures of food, a 50% chance of an increment of sorrow from a negative test, or avoiding the icy plunge) is chosen over a deferred alternative (weight loss, a 50% chance of life free from the threat of Huntington's disease, or cool relief). Moving the consequences of choice away from the present, while holding constant everything else about them, often reverses the preference order. For eating and for taking the plunge, it is plain that the preference reverses. For Huntington's disease, we can surmise that it also reverses, because most of us would advise a person at risk to take the test (as physicians now do advise them), but are likely to be unable to do so when we face the prospect of immediate bad news ourselves.

In each case, the discounting factor applied to restraint in relation to impulse shrinks as it moves further in time, so we choose impulsively when the consequences are at hand, but with restraint when they are deferred. We are disposed to see things in better perspective as they become more remote. How come?

One approach is to invoke a systematic psychophysical distortion of time perception, foreshortening remote time intervals. That may, indeed, be true, but an answer[3] closer to the data and of more fundamental significance is that we discount events hyperbolically in time (at least approximately; Ainslie, 1975; Chung & Herrnstein, 1967; Mazur, 1985, 1987; Williams, 1988), rather than exponentially, as rational choice theory assumes. A hyperbolic time discounting function has, as one of its corollaries, the very foreshortening of remote time inter-

[2] I owe this comparison to George F. Loewenstein.

[3] The answer is contemporary, but the question of time perspective in choice is not. I thank James Q. Wilson for calling my attention to David Hume's characterization of it in the 18th century, from an essay on the origins of government:

> When we consider any objects at a distance, all their minute distinctions vanish, and we always give the preference to whatever is in itself preferable, without considering its situation and circumstances. . . . In reflecting on any action which I am to perform a twelvemonth hence, I always resolve to prefer the greater good, whether at that time it will be more contiguous or remote; nor does any difference in that particular make a difference in my present intentions and resolutions. My distance from the final determination makes all those minute differences vanish, nor am I affected by any thing but the general and more discernible qualities of good and evil. But on my nearer approach, those circumstances which I at first overlooked begin to appear, and have an influence on my conduct and affections. A new inclination to the present good springs up, and makes it difficult for me to adhere to my first purpose and resolution. This natural infirmity I may very much regret, and I may endeavor, by all possible means, to free myself from it. (Hume, 1777/1826, pp. 314–315)

vals that the data suggest. With exponential discounting, the discount rate remains fixed; with hyperbolic, the rate itself shrinks with time.

Exponential time discounting arises from rationalistic considerations; hyperbolic time discounting is a frequent result of behavioral experiments on various species, including human. The evidence for hyperbolic discounting comes primarily from choice experiments in which it is assumed that the subjects are obeying the matching law, a principle of choice that has been widely observed in the laboratory and is defined here in the context of the next riddle to be discussed (Ainslie, 1975; Chung & Herrnstein, 1967; Herrnstein, 1981; Mazur & Herrnstein, 1988).

Imagine that a person is playing tennis, and her or his opponent comes to the net (Herrnstein, 1989; Herrnstein & Mazur, 1987). Assume that the person must now choose between a lob and a passing shot and disregard, for simplicity, any strategic plan in which the opponent may be engaging. Consider the opponent a random variable. Both lobs and passing shots are more effective if they are surprising, and less effective if they are expected. Assume, finally, that surprise has a larger effect on the effectiveness of lobs than of passing shots, which is probably the case. How does he or she decide which shot to hit?

I have presented this riddle to many people, including devotees of rational choice theory. Almost everyone who agrees to play along comes up with something like the following: "As long as one shot is more effective than the other, I'd use it. As I use it, the surprise factor takes its toll. When the other shot becomes more effective, I'd switch to that one. And so I'd oscillate from one shot to the other, trying to switch to the one that is currently more effective."

No one to whom I have presented the riddle has ever spontaneously noticed that the strategy I just characterized may be significantly suboptimal. Some concrete values may help. Suppose the lob has a .9 chance of earning a point when it is a surprise and a .1 chance of doing so when it is fully expected. A surprise passing shot, we can assume, has a .4 chance of being effective, and a .3 chance if it is fully expected. Figure 1 plots these points and connects them linearly for intermediate levels of expectation, as functions of the expectation for a lob. The dashed curve is the joint effect of both shots, which is to say, the average of their effectivenesses weighted by the relative frequency of their use. Figure 1 assumes that expectations for the two shots are determined by the probability of their use in the recent past and that the probability of one is the complement of that for the other.

The strategy that people espouse falls at the intersection of the two solid lines in Figure 1. It is here, at about two thirds lobs, that the two shots have equal effectiveness. A shift toward more lob use reduces the effectiveness of lobs and likewise for more passing shot use. This is a point of equilibrium in the sense that deviations from it are self-negating, if the player is using the strategy of comparing the effectiveness of the shots.

Figure 1
Points Per Shot for a Hypothetical Tennis Player Choosing Between Lobs and Passing Shots as Functions of the Current Probability of Lobs

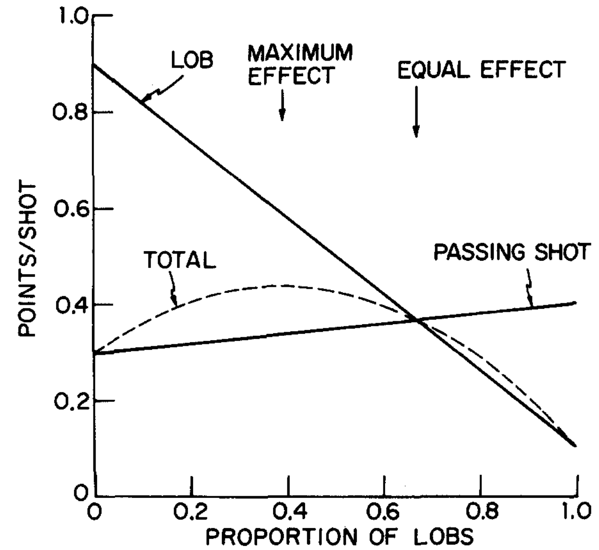

Note. Both shots profit from surprise, but lobs do so more than passing shots. The behavioral equilibrium point is at about two thirds lobs, but the optimal strategy is at about 40% lobs. Data are from "Darwinism and Behaviorism: Parallels and Intersections" by R. J. Herrnstein. In *Evolution and Its Influence* edited by A. Grafen, 1989, London: Oxford University Press. Copyright 1989 by Oxford University Press. Data are also from "Making up Our Minds: A New Model of Economic Behavior" by R. J. Herrnstein and J. E. Mazur, 1987, *The Sciences*, November/December. Copyright 1987 by New York Academy of Sciences. Used by permission.

If the player were a point-maximizer, however, she or he would use a different strategy. The player would look at the two shots overall and pick the point at which their joint effectiveness is at a maximum, shown in Figure 1 as the maximum of the dashed curve, near 40% lobs. At the maximum, each lob is more effective than each passing shot, but the two of them together provide the highest returns. Even after I try to explain where the maximum strategy lies, many people express puzzlement. Finding the maximum in a situation like this does not seem to come naturally.

What does come naturally, as noted earlier, is the strategy that stabilizes at the intersection of the two solid lines, where both shots have the same average value in points. This distribution of shots is dictated by the matching law. According to the matching law, behavior is distributed across alternatives so as to equalize the reinforcements per unit of behavior invested in each alternative. Or to put it another way, the proportion of behavior allocated to each alternative tends to match the proportion of reinforcement received from that alternative. The tennis riddle thus provides an example of spontaneous human irrationality and of the relation of that irrationality to the matching law.

In several hundred experiments, mainly on animals but also on human beings, choice has approximately conformed to the matching law (for recent reviews of the literature, see Davison & McCarthy, 1988; Williams, 1988). The complexities of this literature would be out of place here, but simply stated, the widely accepted conclusion is that subjects allocate behavioral alternatives so that each alternative action earns the same rate of reward per unit of behavior invested, once variations in response topography and in reward quality have been taken into account. This equalization of reward rates is the matching law and is exemplified by the allocation people choose in the tennis riddle. Although the matching law is well established, there are varying explanations for its occurrence.

William Vaughan, Drazen Prelec, and I did an experiment (described in Herrnstein and Prelec, 1989) on human subjects that is reminiscent of the tennis riddle. Its results suggest what the dynamic process is that causes people (and animals) to obey the matching law. Volunteer subjects spent a half hour or so in an experimental booth containing two response keys. The subjects were told they would earn a few cents every time they depressed either key when the trial light was illuminated. The probability of reinforcement was, in other words, 1.0 under all conditions. Each trial was separated from the next by an intertrial interval. The intertrial interval following a choice of one of the keys (the 1-key) was two seconds shorter than that following a choice of the other key (2-key). However, the intertrial interval following either choice was an increasing linear function of the proportion of 1-key choices in the preceding 10 trials (including the present one). Figure 2 lays out the relations.

Figure 2 shows the intertrial interval following 1- and 2-key choices (dashed and solid lines, respectively) as functions of the proportion of 1-key choices during the 10 preceding trials and the weighted average of both delays (dotted line). After the verbal instructions, subjects were given 100 free trials to familiarize themselves with the procedure, then a 10-minute period for playing "for keeps." They were advised that it was in their best interest to complete as many trials as possible in the 10-minute period, but were told nothing more about the contingencies.

The 17 dots show the performances of 17 subjects in a single session. Some subjects worked for extended times over multiple sessions, without any systematic change in their patterns of responding. The optimal strategy would have been to choose key 2 every time, thereby minimizing the intertrial interval at 4 seconds. Instead, all but one of the subjects chose key 1 most of the time, and a few chose it virtually all of the time, enduring the worst possible overall delay between trials, 6 seconds.

Having been subjects ourselves (though not included among the 17 plotted), we think we know what happens. A two-second difference in intertrial intervals is hard to disregard. It therefore feels right to add more key 1 responses into the mix of choices. One strategy that is immune to that temptation is exclusive preference of key 1, which is the worst possible strategy. Subjects sometimes sense that their choices are influencing the intertrial intervals, but they do not know what to make of the information. Likewise, with the tennis riddle, telling people exactly how each shot's effectiveness interacts with its use does not guide them to a proper application of the information. The next section attempts to characterize the underlying process in general terms.

The rational choice advocates I talk to are likely to complain that these two examples—the tennis riddle and the intertrial interval experiment—are too complex for the basic maximizing tendency to emerge. "Give people a chance," they say, "and they will maximize. Of course, if you mix them up badly enough, they won't." That argument calls for two rhetorical rejoinders. One is that if these contingencies are complex enough to suppress our basic maximizing tendency, then it is a fragile tendency indeed. One may argue about how to define simplicity and complexity, but not many decision problems in ordinary life are simpler by any defensible definition of those terms.

The second reply is that it would have been uncharacteristic for natural selection to endow us with a decision rule, such as the maximizing principle, and to fail to endow us with the capacity to exploit the relevant variables. Creatures are usually remarkably sensitive to the stimuli controlling their behavior—like fish responding to the currents in the water around them. With maximization problems, we seem to be like fish out of water.

Figure 2
Delay Between Trials Following Key 1 (Dashed Line) and Key 2 (Solid Line) Choices, as Functions of the Proportion of Key 1 Choices in the Preceding 10 Trials

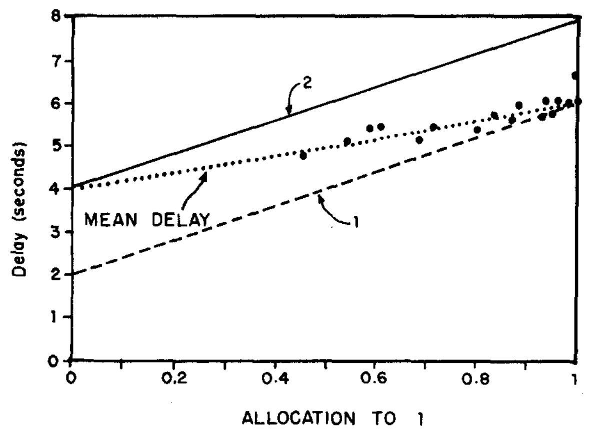

Note. The dotted line is the weighted average of the delays following the two choices, at the prevailing allocation to key 1. A single session's behavior for each of 17 subjects is shown by the dots. Data are from an unpublished experiment by Prelec, Vaughan, and Herrnstein.

Melioration and Matching

More to the point, however, is that the results for both of these examples conform to a familiar finding in the behavior laboratory (Herrnstein, 1982; Herrnstein & Vaughan, 1980; Prelec, 1982), one that may explain the matching law. Confronted with choices that provide differing rates of reward per unit of behavior invested, organisms allocate more time to the alternatives that provide the higher rates of reinforcement. This tendency to shift behavioral allocation toward more lucrative alternatives has been called melioration.

As a subject meliorates, that is, as it shifts toward better alternatives, its doing so may cause the reinforcement returns from the alternatives to change, for various reasons.[4] Behavior then continues to shift toward the newly better alternatives. We would shift from lobs to passing shots and vice versa as one or the other was doing better on the average. In the intertrial experiment, subjects shift toward key 1 because it is followed by a shorter delay. The principle of melioration seems commonsensical enough; the only surprises may be that it fails to maximize reinforcement and that it leads to an equilibrium dictated by the matching law.

The melioration principle says that choice is driven, in effect, by a comparison of the average returns from the alternatives. Equilibrium is attained either when one alternative has displaced all others or all the remaining alternatives from the choice set are providing equal returns per unit consumption. Either of these equilibrium conditions conforms to the matching law, according to which the relative frequency of each behavioral alternative matches the relative frequency of reinforcement provided by it (Herrnstein & Vaughan, 1980). The relation between the matching law and hyperbolic time discounting is exemplified in Chung and Herrnstein's experiment on delayed reinforcement (1967) and in a large variety of procedures explored by Mazur (Mazur 1985, 1987; Mazur, Snyderman, & Coe, 1985; Mazur & Vaughan, 1987; see also the summary in Williams, 1988).

Melioration and matching can produce suboptimalities when one's allocation of choices affects the returns we obtain from the alternatives, as in both the tennis riddle and intertrial experiment (see Herrnstein & Prelec, 1989, for a formal treatment). Such interactions are the rule in everyday life, not the exception. When they take place, the natural decision-making tendency is evidently to disregard the implications of this interaction for overall returns and to focus instead on the current average returns from the alternatives, which is to say, to meliorate, hence to match.

In the tennis riddle, the familiar strategy exemplifies melioration—shift from one shot to the other when the other becomes more effective. In the intertrial experiment, melioration drives choice toward key 1, the alternative with the shorter delay following it, which maximally defies rational choice theory. Comparably stark violations of rational choice theory have been observed in more systematic experiments.

For example, pigeons in an experiment chose between two alternatives, each delivering a bit of food at irregular intervals with average values that were varied parametrically during the course of the study (Heyman & Herrnstein, 1986). For one alternative, the clock that timed the intervals ran only when the pigeon was choosing that alternative; for the other alternative, the clock ran all the time, but the reinforcements it scheduled were only given to the pigeon after it chose that alternative. The pigeon could switch from one alternative to the other at will, with a single peck at a key.

To earn the maximal reinforcement rate, the pigeons should have spent most of their time on the alternative for which the clock ran only when the alternative was chosen, sampling the other alternative occasionally, to collect reinforcements that were due. But maximizing reinforcement here would violate melioration and matching. Numerous workers have lately studied variants of this schedule (known as a concurrent variable-interval, variable-ratio schedule) because it sharply discriminates between the melioration principle and the reinforcement maximizing principle. The evidence to date has clearly, if not unanimously, favored melioration (see reviews in Heyman & Herrnstein, 1986, and Williams, 1988).

A particular transition in one experiment is shown in Figure 3. Midway through the experiment, two pigeons had just started working on a new pair of schedule values. The reinforcement maximization principle predicted that they would spend none of their time on the alternative called VT, the one for which the clock runs continuously. The melioration principle predicted that they would spend almost all their time choosing VT. Because of the schedule parameters in the preceding condition, the pigeons started off spending a quarter or less of their time on VT, as the open points indicate. They started off at a reinforcement rate close to the maximum possible.

However, as the dashed arrows indicate, they also started off in this new condition quite far from the point predicted by the matching law. Over several weeks of daily sessions, the pigeons gradually shifted in their choices toward the VT alternative, as melioration drove them toward conformity with the matching law. The dashed arrows shrunk to virtually nothing, indicating conformity with matching. As they did, earnings fell almost steadily, traced by the triangular-shaped points. Obeying the matching law cost the pigeons more than a third of their over-all rate of food reinforcement. This condition's results were typical. In the experiment as a whole, the pigeons earned food at a lower rate than they would have by allocating choices randomly to the two alternatives, let alone what they could have earned as food reinforcement maximizers.

Does real life ever arrange anything like this schedule for pigeons or other species? Houston (1986) has shown that a schedule much like it governs food availability for the pied wagtail, a bird living in the Oxford University environs where he and his colleagues work. Field studies confirm that the foraging of the pied wagtail is suboptimal in the way that melioration predicts for this sort of schedule, with too much time being invested in the alternative

[4] They may change for motivational or situational reasons. For example, when a food source is chosen more often, hunger may be reduced, making a quantity of food less reinforcing, or the source may be depleted, reducing the average return per unit of behavior invested.

Figure 3
Two Pigeons on a Version of a Concurrent Variable-Interval, Variable-Ratio Schedule of Reinforcement

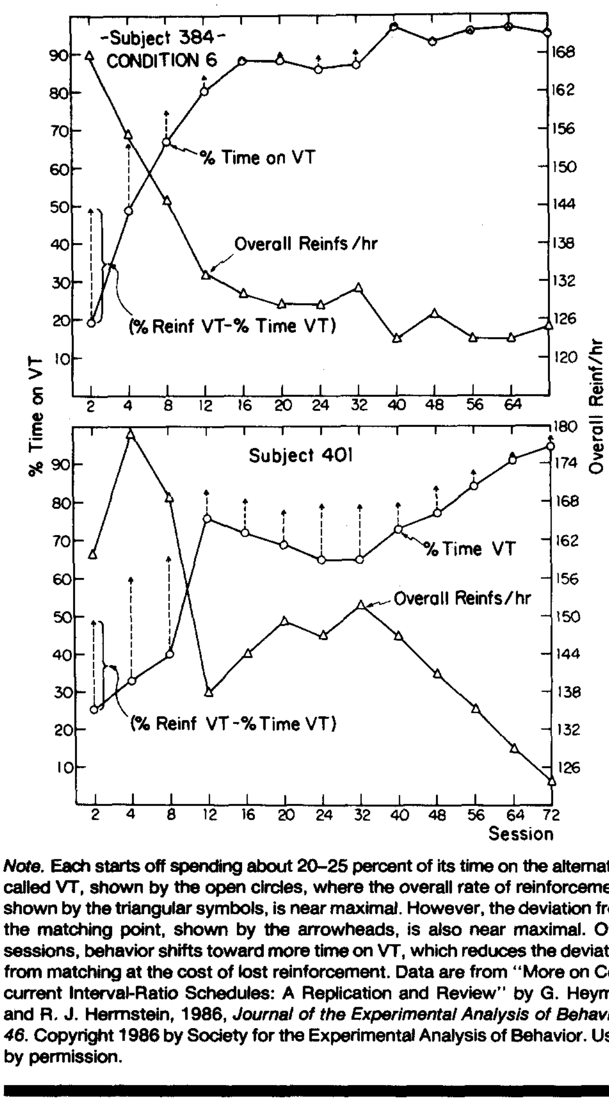

Note. Each starts off spending about 20–25 percent of its time on the alternative called VT, shown by the open circles, where the overall rate of reinforcement, shown by the triangular symbols, is near maximal. However, the deviation from the matching point, shown by the arrowheads, is also near maximal. Over sessions, behavior shifts toward more time on VT, which reduces the deviation from matching at the cost of lost reinforcement. Data are from "More on Concurrent Interval-Ratio Schedules: A Replication and Review" by G. Heyman and R. J. Herrnstein, 1986, *Journal of the Experimental Analysis of Behavior, 46*. Copyright 1986 by Society for the Experimental Analysis of Behavior. Used by permission.

that needs to be sampled only occasionally to gain maximal food recovery.

The melioration principle implies suboptimality particularly for the class of situations that Prelec and I have characterized as "distributed choice" (Herrnstein & Prelec, 1989). For distributed choices, the organism does not make a once-and-for-all decision about the alternatives in a choice set. Instead, repeated choices are made over some period of time, and the decision variable is the allocation among alternatives, the ongoing proportion of lobs versus passing shots, for example.

Most "style of life" questions concern distributed choice, Prelec and I have suggested. At no moment in life does one choose, for example, to become promiscuous or a glutton or an alcoholic. Rather, those human frailties creep up on us insidiously, the result of numerous minor choices, many of which may be barely, if at all, blameworthy. Other examples best considered as matters of distributed choice are easy to think of: the continuum from miserliness to spendthriftiness, from profligacy to prudery, or from being an exercise junkie to being sedentary. No single choice is involved in being diligent rather than lazy, honest rather than dishonest, loyal rather than disloyal, clean rather than dirty, and so on.

Rational choice theory has not formulated a clear or effective framework for distributed choice, although some theorists have addressed particular examples, such as drug addiction. The hope has been that the concepts applied to the idealized, timeless choices in their theory will smoothly generalize to temporally extended fields, to distributed choice. Evidently, they have not. In contrast, in the experimental analysis of behavior the study of behavioral allocation within a period of observation has been central for more than a generation. Both melioration and matching describe behavior in the framework of distributed choice.

Figure 4, for example, is our account of the suboptimality observed in an experiment such as that just summarized. This is the same sort of chart presented earlier for the intertrial interval experiment. The x-axis plots the proportion of time spent on the alternative that needed to be sampled only occasionally because its clock ran continuously (variable interval or VI). As a function of this variable, the y-axis gives the rate of reinforcement received from this alternative while the subject is choosing it. Because its clock runs all the time, the less time the

Figure 4
Schematic Representation of the Reinforcement Structure of a Concurrent Variable-Interval, Variable-Ratio Schedule

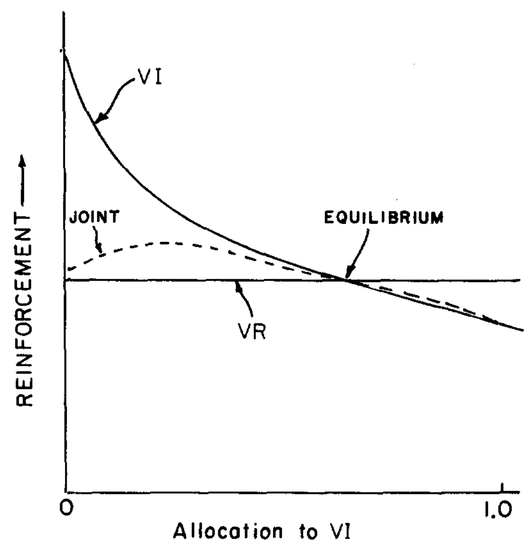

Note. The reinforcement rate for the variable interval (VI) is inversely related to the time allocated to VI; for the variable ratio (VR), reinforcement rate is independent of time allocated to either alternative. Consequently, the equilibrium point (i.e., matching) is displaced from the allocation that would maximize reinforcement.

subject spends on it, the higher the rate of reinforcement when it is sampled. It can be thought of as a model of a reinforcement source that gets depleted as it is sampled and restores itself when it is left unsampled, or as a motivational state that fluctuates with deprivation and satiation.

Also shown in Figure 4, as a function of time spent on this alternative, is the rate of reinforcement received from the other alternative (variable ratio or VR). Because the other clock runs only when the other alternative is being sampled, it provides a fixed rate of return per unit of time invested in it, however much the subject samples it. It can be thought of as a model of a standard gambling device, with a fixed probability of winning. Finally, Figure 4 shows, as a dashed curve, the joint returns from the two alternatives.

When the subject is spending only a small proportion of its time on the VI, then the schedule provides a higher rate of return than the VR. Melioration then dictates increasing time on the VI. However, if the subject spends too much time there, the rate of return falls below that provided by the VR. Melioration then says: "too much VI." Between the extremes falls the equilibrium point, where the alternatives provide equal rates of return per investment. This point conforms to the matching law, as do all equilibrium points produced by melioration.

To maximize reinforcement, the subject would have to find the highest point on the dashed curve. It is generally the case that at the maximum, the subject would be earning a higher rate of return from the VI than from the VR. To maximize, the subject therefore must resist its tendency to spend more time at a more lucrative alternative. Think of it this way: At the maximizing allocation, whatever would be gained by spending more time on the currently more lucrative variable interval is more than lost by the declining rate of return from that schedule. By now, several scores of experimental pigeons, rats, pied wagtails, and perhaps other species, have failed to maximize and have been drawn instead toward the matching point, which always means too much time spent on the VI.

No parametric, systematic study of human subjects earning reinforcement on the concurrent variable-interval, variable-ratio schedule has yet been published, partly because such experiments take months to complete. Anecdotal evidence suggests that the results might not differ fundamentally.

Suppose, for example, that you eat only caviar or hamburger for supper every night, that sometimes you eat one and sometimes the other, and that price is no object. If you chose caviar whenever you thought it would be even marginally better tasting, you would be meliorating and would be allocating your choices so that caviar tasted no better than hamburger on the average, and vice versa. Assuming that the pleasure of caviar declines more with consumption than that of hamburger, your aggregate pleasure could be far less than it is for those of us for whom caviar must remain a special treat. To get the maximum pleasure, we should husband caviar as a treat, rather than equate it to hamburger. If you ate more caviar than is optimal for maximum pleasure, you would be behaving like pigeons and rats on concurrent variable-interval, variable-ratio schedules, spending too much time on the variable interval.

More generally, most, if not all, of us imagine we would be happier wealthier than we are. We feel this way however wealthy we are. Are the very wealthy, then, very happy? Charles Murray (1988) concluded in a new book that neither survey results nor common experience bears out the expectation of continuously increasing happiness with increasing wealth. Not only does more money not necessarily buy more happiness, it often seems to buy trouble. The "poor little rich girl," for example, does exist. Traditional morality counsels against the unfettered pursuit of material wealth but provides no scientific justification for its wise counsel. The melioration principle explains how having more money may lead to counterproductive equilibrium points; the maximization principle does not, at least not without quite a few epicycles.

Addictions and Other Pathological Choices

Figure 4 illustrates distributed choice, given just two alternatives. Visualizations of larger choice sets are harder to concoct, but the evidence suggests that matching and melioration apply as well to larger choice sets, albeit with greater formal complexity (Herrnstein & Prelec, 1989). However, for purposes of explication, we can always collapse larger choice sets to a choice between a particular alternative and "everything else." That is what Prelec and I have recently done in a theoretical discussion of addiction seen as a pathology of distributed choice (Herrnstein & Prelec, 1989).

Figure 5
A Model of Behavioral Addiction

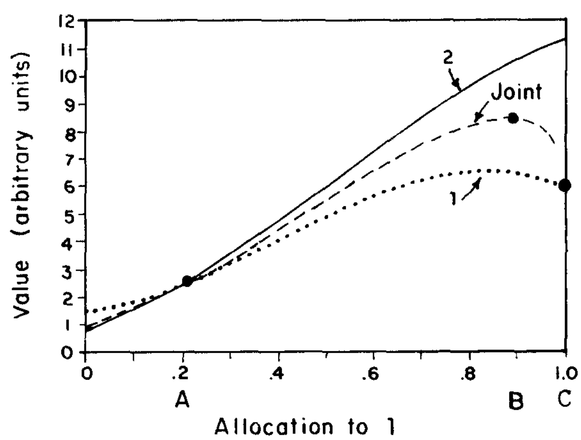

Note. The average value of each of two competing commodities, labeled 1 and 2, declines as allocation shifts from 1 to 2. The joint returns from both commodities, the dashed curve, pass through its maximum at the filled point above B; equilibrium is that at A, and the filled point above C represents the joint returns when allocation is totally to 1. Adapted from *Melioration: A Theory of Distributed Choice* by R. J. Herrnstein and D. Prelec, 1989, Cambridge, MA: Harvard Business School. Copyright 1989 by Harvard Business School. Adapted by permission.

Figure 5 shows how the reinforcement rates, here called "value," of two alternatives, 1 and 2, depend on allocation to alternative 1, where the total allocation, set at 1.0, is divided between the two alternatives. The addictive commodity, alternative 2, shown by the solid curve, declines in value per unit the larger the investment in it. The alternative competing with it, 1, shown by the dotted curve, rises in value the more it is chosen, up to some high level of consumption, beyond which it too falls with subsequent consumption. The joint returns from the two alternatives are plotted by the dashed curve.

Melioration predicts equilibrium at A, where choosing the addictive commodity, 2, predominates. A maximizer should choose the highest point on the joint function, B, but at B, the addictive commodity seems far more reinforcing than its competitors and is therefore, by melioration, likely to be chosen more. Exclusive choice of the nonaddictive alternative, 1, is labeled C.

This structure of reinforcement contingencies, Prelec and I suggest, portrays addiction inasmuch as the reinforcement returns at C are higher than those at A. In other words, in this environment, the subject benefits simply by being deprived of the option of choosing 2. Rational choice theory has trouble dealing with the idea that a person benefits by being deprived of an option, but common experience, and melioration, says it can happen easily.

Figure 5 shows addiction behaviorally, and that may be its special strength because addiction is fundamentally a pathology of behavior. An addictive commodity steeply loses average value per unit consumed at higher levels of allocation to it. At the same time, the alternatives to addiction also lose value when the addictive commodity predominates. As a result, the equilibrium point is at a low overall level of reinforcement, which is why addictions are considered counterproductive. The person may well know how counterproductive his or her way of life is but be helpless to escape from it.

A reinforcement structure containing the essentials of Figure 5 could arise in many ways. If the addictive commodity is a chemical, then excessive use may induce chemical changes in the brain, which in turn induce a tolerance that reduces the hedonic punch of a unit of the addictive chemical and also spoils the pleasures derived from other activities. But the commodity need not be a chemical. People who overwork may find that work is less fun than it used to be, but that it is still better than not working. We may then say the person is a workaholic, implying irrational fixation on work. People sometimes find themselves trapped in bad personal relationships: The keen pleasure in a romance may be long gone, but alternative personal relationships have lapsed and are not easily reconstructed. We may not say that the person is addicted to his or her partner, but we could.

Treating addiction often consists of a total prohibition against the addictive commodity, a shift to C in Figure 5. As long as the person never indulges in the addictive commodity, he or she will not experience its superior returns in this region of allocations. Overall reinforcement at C is suboptimal but superior to that at A. The optimal allocation, at B, permits low levels of indulgence at which the addictive commodity is still more pleasurable than its alternatives. Most people would find B too hard to sustain, but that only shows they are meliorizers, not maximizers.

Another approach to treatment is suggested by the analytic framework. A lowering of the value curve for the addictive commodity shifts the equilibrium point toward the optimal allocation, as in Figure 6. The dashed curve, labeled 2′, illustrates how the addictive commodity curve would change if a use tax is applied to it. The new equilibrium, at A′, is a shift away from excessive indulgence, toward optimality but not reaching it. A comparable improvement would result if the nonaddictive alternative earned a consumption bonus. Raising the height of the dotted curve likewise shifts equilibrium to the right. This tells us nothing more than we know from common sense. But conforming to common sense, with no loss of rigor, is one of the advantages of the melioration framework.

Conclusions

Rational choice theory and the melioration principle converge under certain conditions. For example, when the reinforcement rates associated with competing alternatives are independent of the frequency with which they are sampled, the subject maximizes by choosing the better alternative exclusively. We all choose $10 over $5. But melioration similarly implies exclusive choice of the better alternative here. In the laboratory, this is approximated by concurrent ratio schedules; in economic theory, it is choice among constant probability alternatives, which is the classical economic paradigm. When human or animal subjects maximize in situations like this, as they often

Figure 6
The Equilibrium, at A, in Figure 5 Shifts to A′ When the Value of Commodity 2 is Reduced for All Allocations

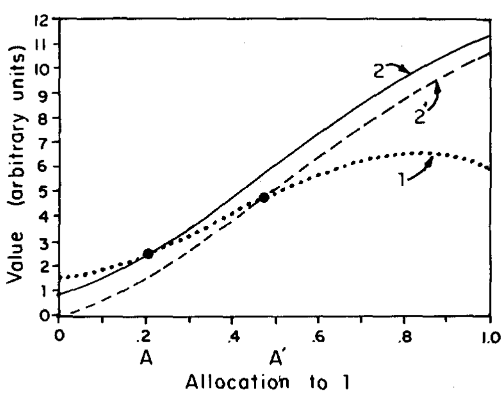

Note. The value of commodity 2 is shown by the dashed curve. Consumption of 2 at equilibrium would be reduced by a consumption tax on 2. Adapted from *Melioration: A Theory of Distributed Choice* by R. J. Herrnstein and D. Prelec, 1989, Cambridge, MA: Harvard Business School. Copyright 1989 by Harvard Business School. Adapted by permission.

do, neither theory is strengthened at the expense of the other.

Rational choice theory adequately describes distributed choice in those situations in which the distributed nature of the choice is immaterial in the sense that the returns do not depend on the frequency of sampling. In many, though not all, other situations, it fares less well. Nothing in rational choice theory can tell us when it fares well and when it fails, but melioration does tell us. When melioration implies utility maximization (i.e., maximum reinforcement rate), rational choice theory adequately describes behavior.

But describing behavior does not seem to be the proper use for rational choice theory. Rational choice theory tells us how choice *should* be allocated, given a reinforcement or utility structure, not how it will be allocated. This normative function it serves admirably and usefully. A better analogy for rational choice theory than Newtonian physics is Boolean algebra.

George Boole, a 19th-century English mathematician, wrote his book (1854/1911) on the binary arithmetic named after him as a description of human reasoning: He titled it *An Investigation of the Laws of Thought.* It turned out not to be much of a theory of human thought, but as a calculus of reasoning, the book was epochal. Instead of discovering the laws of thought, Boole had invented the algebra embodied in all modern digital computers, an algebra uniquely well suited to the behavior of a network of binary switching elements. One reason we find computers helpful is that our thought processes are often not Boolean. In just that way, we need rational choice theory because, as meliorizers, we often act suboptimally. How a meliorizer can make use of the guidance provided by rational choice theory is a complex matter, far from well worked out and beyond the scope of this essay (some sketchy notions are presented in Herrnstein, 1988, 1989).

Rational choice theory lies at the heart of not only modern microeconomic theory but also political doctrines that advocate minimal government—libertarianism and anarchism, for example. The idea is that, insofar as people behave rationally, they should be left to their own devices, except when collective behavior undermines individual interest, as when maximizing fishers overfish the waters or each individual decides that someone else should do a particular job, like serve in the army or build a road. But suppose people fundamentally and individually misbehave, as the evidence indicates they do. Then we would expect government to take account, not just of the defects of collective action, but of individual action as well, as David Hume (1777/1826) said more than 200 hundred years ago. As old as it is, the idea remains unexplored and revolutionary, and it defines a conceptual frontier that students of the experimental analysis of behavior are uniquely well qualified to cross.

REFERENCES

Ainslie, G. (1975). Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin, 82,* 463–496.

Boole, G. (1854). *An investigation of the laws of thought, on which are founded the mathematical theories of logic and probabilities.* London: Walton and Maberly.

Brody, J. E. (1988, August 25). Personal health. *New York Times,* p. 317.

Chung, S. H., & Herrnstein, R. J. (1967). Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior, 10,* 67–74.

Davison, M., & McCarthy, D. (1988). *The matching law: A research review.* Hillsdale, NJ: Erlbaum.

Dawes, R. M. (1988). *Rational choice in an uncertain world.* San Diego, CA: Harcourt Brace Jovanovich.

Frank, R. H. (1988). *Passions within reason: The strategic role of the emotions.* New York: Norton.

Heiner, R. A. (1983). The origin of predictable behavior. *American Economic Review, 73,* 560–595.

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior, 13,* 243–266.

Herrnstein, R. J. (1981). Self control as response strength. In C. M. Bradshaw, E. Szabadi, & C. F. Lowe (Eds.), *Quantification of steady-state operant behaviour* (pp. 3–20). Amsterdam: Elsevier/North Holland Biomedical Press.

Herrnstein, R. J. (1982). Melioration as behavioral dynamism. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of Behavior, Vol. 2: Matching and maximizing accounts* (pp. 433–458). Cambridge, MA: Ballinger.

Herrnstein, R. J. (1988). A behavioural alternative to utility maximization. In S. Maital (Ed.), *Applied behavioral economics* (pp. 3–60). London: Wheatsheaf Books.

Herrnstein, R. J. (1989). Darwinism and behaviorism: Parallels and intersections. In A. Grafen (Ed.), *Evolution and its influence* (pp. 35–61). London: Oxford University Press.

Herrnstein, R. J., & Mazur, J. E. (1987, November/December). Making up our minds: A new model of economic behavior. *The Sciences,* pp. 40–47.

Herrnstein, R. J., & Prelec, D. (1989). *Melioration: A theory of distributed choice* (Harvard Business School Working Paper, 89-030). Cambridge, MA: Harvard Business School.

Herrnstein, R. J., & Vaughan, W., Jr. (1980). Melioration and behavioral allocation. In J. E. R. Staddon (Ed.), *Limits to action: the allocation of individual behavior* (pp. 143–176). New York: Academic Press.

Heyman, G., & Herrnstein, R. J. (1986). More on concurrent interval-ratio schedules: A replication and review. *Journal of the Experimental Analysis of Behavior, 46,* 331–351.

Hirshleifer, J. (1982). Evolutionary models in economics and law: Cooperation versus conflict strategies. *Research in law and economics, 4,* 1–60.

Houston, A. I. (1986). The matching law applies to wagtails foraging in the wild. *Journal of the Experimental Analysis of Behavior, 45,* 15–18.

Houston, A. I., & McNamara, J. M. (1988). A framework for the functional analysis of behaviour. *Behavioral and Brain Sciences, 11,* 117–163.

Hume, D. (1826). *The philosophical works* (Vol. 2). Edinburgh: Black and Tait. (Original work published 1777)

Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases.* Cambridge, England: Cambridge University Press.

Luce, R. D. (1988). Rank-dependent, subjective expected-utility representations. *Journal of Risk and Uncertainty, 1,* 305–332.

Luce, R. D. (1989). *Linear utility models with rank- and sign-dependent weights.* Unpublished manuscript.

Machina, M. J. (1987). Decision-making in the presence of risk. *Science, 236,* 537–543.

Margolis, H. (1987). *Patterns, thinking, and cognition: A theory of judgment.* Chicago: University of Chicago Press.

Mazur, J. E. (1985). Probability and delay of reinforcement as factors in discrete-trial choice. *Journal of the Experimental Analysis of Behavior, 43,* 341–351.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behavior, Vol. 5: The effect of delay and of intervening events on reinforcement value* (pp. 55–73). Hillsdale, NJ: Erlbaum.

Mazur, J. E., & Herrnstein, R. J. (1988). On the functions relating delay, reinforcer value, and behavior. *Behavioral and Brain Sciences, 11,* 690–691.

Mazur, J. E., Snyderman, M., & Coe, D. (1985). Influences of delay and rate of reinforcement on discrete-trial choice. *Journal of Experimental Psychology: Animal Behavior Processes, 11,* 565–575.

Mazur, J. E., & Vaughan, W., Jr. (1987). Molar optimization versus delayed reinforcement as explanations of choice between fixed-ratio and progressive-ratio schedules. *Journal of the Experimental Analysis of Behavior, 48,* 251–261.

Murray, C. (1988). *In pursuit of happiness and good government.* New York: Simon & Schuster.

Prelec, D. (1982). Matching, maximizing, and the hyperbolic reinforcement feedback function. *Psychological Review, 89,* 189–230.

Stigler, G. J., & Becker, G. S. (1977). De gustibus non est disputandum. *American Economic Review, 67,* 76–90.

Vaughan, W., Jr., & Herrnstein, R. J. (1987). Stability, melioration, and natural selection. In L. Green & J. H. Kagel (Eds.), *Advances in behavioral economics* (Vol. 1, pp. 185–215). Norwood, NJ: Ablex.

Williams, B. A. (1988). Reinforcement, choice, and response strength. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Eds.), *Stevens' handbook of experimental psychology, Vol. 2* (pp. 167–244). New York: Wiley.