1-2003

Preferences and Rational Choice: Introduction

Claire Oakes Finkelstein
University of Pennsylvania Carey Law School

Matthew D. Adler
University of Pennsylvania Carey Law School

Peter H. Huang
University of Pennsylvania Carey Law School

Repository Citation

SYMPOSIUM

PREFERENCES AND RATIONAL CHOICE:
NEW PERSPECTIVES AND LEGAL IMPLICATIONS

INTRODUCTION

MATTHEW D. ADLER,[†] CLAIRE FINKELSTEIN,[††] & PETER H. HUANG[†††]

The fifteen articles in this issue were originally written for a Symposium entitled *Preferences and Rational Choice: New Perspectives and Legal Implications*, held at the University of Pennsylvania Law School on March 1-2, 2002. The articles focus on the traditional economic account of individual rationality and the implications of recent criticisms of that account for the law. The criticisms in many cases arose by applying the insights and methods of other disciplines to this fundamental problem in the economic literature. The issue brings together economists, philosophers, psychologists, business and finance schol-

[†] Professor of Law, University of Pennsylvania Law School.
[††] Professor of Law and Philosophy, University of Pennsylvania Law School.
[†††] Assistant Professor of Law, University of Pennsylvania Law School and Visiting Professor of Law, University of Virginia School of Law.

(707)

ars, and lawyers in an effort to shed light on a question crucial for all these disciplines: what is the correct account of individual human rationality? In this Introduction, we will sketch the contribution the articles collected here make to the ongoing debate on this topic and explain the implications of that debate for the law.

The traditional economic account treats rational agents as seeking to maximize their preferences. The preference-maximization model has long been the dominant approach not only in economics, but also in related fields such as psychology, decision theory, philosophical rational choice, and law and economics. The economic model understands the notion of preference in terms of an agent's ranking of her choices and of the various possible outcomes of those choices. In the simplest case, where the agent knows for certain which outcome will result from each choice, she is enjoined to choose the option leading to the highest-ranked outcome. From a philosophical perspective, the above criterion may seem to leave the notion of "preference" an empty one, since it provides no way to assess the rationality of an agent's choices with respect to some underlying state of satisfaction, and so no way to assess independently whether an agent's choices maximize her preferences. The traditional model does not restrict the *content* of an agent's preferences. There is nothing irrational, for example, about choosing to spend the day picking blades of grass instead of making money or writing articles.

But the traditional account does place certain *formal* constraints on the notion of a preference, and these help to make the preference-maximization criterion more robust. In particular, preferences must conform to the following three criteria: (1) *completeness*—an agent must be able to rank any two items with which she is presented, unless she is indifferent between the two; (2) *transitivity*—if an agent would prefer an apple to an orange, and an orange to a banana, then it must be the case that she would prefer an apple over a banana; and (3) *reflexivity*—an agent must be indifferent between an item and an identical item. An agent whose choices do not conform to these conditions would be thought irrational, and her preferences could not be coherently maximized.

What if an agent is choosing under circumstances of incomplete knowledge? How can an agent's choices conform to the above criteria if she does not know, for example, what the results of her choices will be? The traditional model assumes that even if the agent does not know with certainty the outcomes of her choices, she can evaluate the options open to her if she at least knows what the chances are of end-

ing up with one result or another. The model also assumes that the agent's preferences can be measured numerically in terms of utility. The model then stipulates that the utility of a given choice is equal to the utility to the agent of one of the possible outcomes from that choice, discounted by the probability of its occurrence, plus the utility of the next possible outcome from that choice, discounted by its probability of occurrence, and so on for all possible outcomes of a given choice. The best choice for a rational agent, then, is the choice that maximizes her preferences or "utility," adjusting for the probabilities of which she is aware of obtaining specific outcomes.

What intellectual function is performed by the traditional economic account of rationality, and more generally by any model of rational choice? In his article for the symposium issue, Lewis Kornhauser carefully addresses this question, showing how models of rational choice are employed to describe, explain, and evaluate choices, and to facilitate the design of institutions.[1] Much work in experimental psychology, economics, and recently in law addresses the first of these two issues. The crucial question is to what extent the traditional model accurately describes and predicts individual choices. So-called behavioral economists say that the answer is, "Not very well."

Some of the articles make important contributions to this debate about the descriptive and predictive accuracy of the preference-maximization model. Jason Johnston considers various ways in which the behavior of human beings seems to be contradictory. For example, people routinely spend large amounts of money to protect themselves against risks due to carcinogens and toxins, at the same time that they engage in risky behavior like overeating or skydiving. But Johnston argues that this behavior is actually consistent with the traditional model, and is not necessarily the result of irrationally placing too much weight on low-probability events.[2] Similarly, Michael Wachter suggests that the failure of market prices to reflect the true value of a corporation might reflect not the irrational enthusiasm or pessimism of investors, but rather information asymmetries between managers and investors concerning the amount and variability of the future cash flows generated by corporate projects.[3]

---

[1] Lewis A. Kornhauser, *The Domain of Preference*, 151 U. PA. L. REV. 717 (2003).

[2] Jason Scott Johnston, *Paradoxes of the Safe Society: A Rational Actor Approach to the Reconceptualization of Risk and the Reformation of Risk Regulation*, 151 U. PA. L. REV. 747 (2003).

[3] Michael L. Wachter, *Takeover Defense When Financial Markets Are (Only) Relatively Efficient*, 151 U. PA. L. REV. 787 (2003).

A different ongoing debate concerns the success of the prefer-
ence-maximization model as a *normative* account of choice. Sympo-
sium contributors focus on three different ways in which the tradi-
tional model might be questioned. First, a number of articles discuss
the impact of *the passage of time* on an agent's choices. George Ainslie
and John Monterosso agree with the traditional model that agents are
rational to discount future goods.[4] But Ainslie and Monterosso also
argue that a commonly observed form of discounting, known as "hy-
perbolic" discounting, in which the value of a good is inversely pro-
portional to delay, amounts to a kind of irrationality. Hyperbolic dis-
counting, according to Ainslie and Monterosso, "lead[s] to changes of
relative valuation among goods at different delays as time passes."[5] An
agent, for example, may prefer not binging tomorrow to feeling
bloated the next day, even though tomorrow he will prefer to binge.

Note that hyperbolic discounting is consistent with the traditional
model. The agent at each time has a complete ranking of outcomes
which his choice at that time maximizes. Although this ranking varies
intertemporally, the discounting formula that produces it does not.
Indeed, as Leo Katz shows, the phenomenon of intertemporal variabil-
ity in preferences over outcomes is paralleled by similar phenomena
in the moral domain, namely intertemporal variability in judgments of
harm and of wrongfulness.[6] Our ex ante and ex post judgments of the
harmfulness of an action might well diverge. One example Katz gives
is of an assassin who has just fired a bullet at his target but who has
missed.[7] Imagine he is about to fire a second bullet. Before he fires,
one might consider the wickedness of this second attempt on his vic-
tim's life. At that point, we would have to conclude that it is every bit
as bad as the first attempt. Yet once the assassin actually fires, assum-
ing he misses once again, it seems unlikely that what he has done in
firing two bullets, and missing, is appreciably worse than firing one.

Joe Mintoff challenges the traditional model by arguing that ac-
tions that fail to maximize an agent's preferences or utility may be ra-
tional as long as they are selected in the context of a prior intention or

---

[4] George Ainslie & John Monterosso, *Will as Intertemporal Bargaining: Implications for Rationality*, 151 U. PA. L. REV. 825 (2003).

[5] *Id.* at 830.

[6] Leo Katz, *Before and After: Temporal Anomalies in Legal Doctrine*, 151 U. PA. L. REV. 863 (2003).

[7] *Id.* at 863.

plan.[8] He says that a rational agent may adopt plans that call for a sub-optimal action when comparing the cost of complying with the plan to the net benefit the plan provides. If the net benefit exceeds the cost, the plan is rational for him to adopt. Mintoff then argues that the intention an agent forms acts as a filter on the agent's future deliberations. So despite the fact that an individual might fare better by performing an action that is inconsistent with his intention to perform the suboptimal action, it is rational for the agent to stick to his intention when that intention is an essential part of an optimal plan.

A second way in which the traditional model arguably provides a normatively flawed account of choice concerns *risk*. The traditional model allows that an agent can be risk averse or risk loving with respect to the components of outcomes, for example, the agent's wealth. An agent might prefer the certainty of one hundred dollars over a ten percent chance of one thousand dollars. But this model insists that agents are necessarily risk neutral with respect to *utility* itself. Ned McClennen criticizes this feature of the traditional model, suggesting specifically that agents rationally seek to avoid losses with respect to "primary goods"—goods such as bodily integrity, mental and emotional integrity, and all-purpose resources such as money.[9] McClennen seems to suggest that agents might particularly dislike taking risks with primary goods, even when the amounts of those goods are calibrated in utility terms, and not merely when calibrated in some nonutility index.

Claire Finkelstein makes a related point in her article, arguing that being at risk can itself be a kind of harm.[10] Intuitively, for example, subjecting someone to a substantial risk of physical harm seems itself to be a welfare setback—or, to put the point in preference terms, an intrinsically dispreferable feature of outcomes. Thus an agent who is exposed to a risk of harm, but who suffers no actual harm would be in a worse position than another individual who also suffered no harm, but who was not exposed to risk. Similarly, she argues that a person who is exposed to a chance of benefit receives a benefit because of the exposure, even if she does not in the end gain the actual benefit for which she had hoped. Both claims appear to violate the assumptions of the traditional model, because the model assumes that

---

[8] Joe Mintoff, *Can Utilitarianism Justify Legal Rights with Moral Force?*, 151 U. PA. L. REV. 887 (2003).

[9] Edward F. McClennen, *Prudence and Constitutional Rights*, 151 U. PA. L. REV. 917 (2003).

[10] Claire Finkelstein, *Is Risk a Harm?*, 151 U. PA. L. REV. 963 (2003).

an agent's preferences are entirely a function of the actual outcomes she can expect.

Cass Sunstein also engages the subject of risk in his article on the precautionary principle.[11] Sunstein does not commit himself to the traditional model, but he does provide strong arguments against the precautionary principle, which he takes to be a competitor to the traditional model. The precautionary principle is widely espoused by environmentalists and has been endorsed by the United Nations and many foreign governments, particularly in Europe.[12] In its strong form, it says something like the following: "'When an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause-and-effect relationships are not fully established scientifically.'"[13] Sunstein argues that this is incoherent, since precautionary measures will themselves, quite typically, create risks of health or environmental harms.

A final controversial aspect of the traditional model, taken as a normative account, concerns the permissible *content of agents' preferences*. Peter H. Huang, Eric Posner, Jonathan Baron, and Bruce Chapman all, in different ways, address this aspect of the traditional model. Huang demonstrates how the traditional model of preferences can incorporate agents' feeling such emotions as guilt or pride—and not merely be about preferences over their wealth or consumption.[14] Huang demonstrates that (fiduciary) law is able to create a socially desirable equilibrium by influencing individual investors' expectations about how broker-dealers will behave, and those investors' expectations in turn influence the psychological payoffs to broker-dealers. This is a more nuanced, and hence more credible, account of why fiduciary law matters than accounts that explain fiduciary law in terms of the preferences or constraints of broker-dealers. Posner makes the reciprocal point: the traditional model permits agents to have "greedy" preferences, preferences that are wholly focused on the agent's wealth or consumption and ignore all other (intuitively) preferable features of outcomes, such as emotions, relationships, fair

---

[11] Cass Sunstein, *Beyond the Precautionary Principle*, 151 U. PA. L. REV. 1003 (2003).

[12] For a survey of the development and application of the precautionary principle, see THE PRECAUTIONARY PRINCIPLE AND INTERNATIONAL LAW: THE CHALLENGE OF IMPLEMENTATION (David Freestone & Ellen Hey eds., 1996).

[13] Sunstein, *supra* note 11, at 1006 (quoting *Lessons from Wingspread*, in PROTECTING PUBLIC HEALTH AND THE ENVIRONMENT: IMPLEMENTING THE PRECAUTIONARY PRINCIPLE app. A at 353-54 (Carolyn Raffensperger & Joel A. Tickner eds., 1999)).

[14] Peter H. Huang, *Trust, Guilt, and Securities Regulation*, 151 U. PA. L. REV. 1059 (2003).

dealing, accomplishment, and so on.[15]   Posner shows how judges sometimes take greedy preferences to be problematic, even irrational, but does not himself endorse the irrationality claim—in part given the difficulty of distinguishing between appropriately "moderate" and inappropriately "greedy" preferences for dollars and goods.

Baron observes that the traditional model allows persons to have what he calls "moralistic" goals (or preferences): "goals for the behavior of others that are independent of the others' goals."[16] Moralistic preferences are, if not irrational, then morally unappealing, for we would all generally be better off if no one had them.

Finally, Bruce Chapman points to a way in which the traditional model may be too stringent with respect to the substantive content of an agent's preferences.[17] Consider the case of an agent who, when offered a large apple and a large orange, chooses the large apple. But when offered the choice among both these fruits plus a smaller apple, he chooses the large orange, for the reason that it would be a breach of etiquette to choose the bigger of two items in the presence of the smaller. This choice behavior violates the so-called Weak Axiom of Revealed Preference (WARP), which is often seen as an implication of the traditional model: "if an agent ever chooses an alternative x over alternative y from some set of alternatives, then that agent should never (on pain of inconsistency) choose alternative y over alternative x from any other set of available alternatives."[18] In effect, WARP precludes the context of choice from counting as a preferable or dispreferable aspect of the outcomes resulting from choice. Chapman argues, however, that there are instances in which it is rational for context to matter to an individual decision maker.

So much for the possible deficiencies of the traditional model as a model of rationality. Why does any of this matter for the law? Clearly the success or failure of a particular explanation of human behavior has legal implications: which decisions people can be expected to make is surely of great importance to judges, legislators, and administrators alike. But why does the success or failure of a given *normative* account of rationality have relevance for the law? For example, why would it matter to policymakers whether ideally rational agents would

---

[15] Eric A. Posner, *The Jurisprudence of Greed*, 151 U. PA. L. REV. 1097 (2003).

[16] Jonathan Baron, *Value Analysis of Political Behavior—Self-Interested : Moralistic :: Altruistic : Moral*, 151 U. PA. L. REV. 1135 (2003).

[17] Bruce Chapman, *Rational Choice and Categorical Reason*, 151 U. PA. L. REV. 1169 (2003).

[18] *Id.* at 1175.

adopt Mintoffian plan-rationality or McClennenesque prudence? The symposium authors address this important question in a variety of ways. The following are some of the more specific suggestions on this topic.

First, the state might have reason to discourage citizens from making irrational choices and encourage them to make rational ones. This would, arguably, make sense if the action that promoted a person's welfare were also the action it was rational for him to select. Colin Camerer, Sam Issacharoff, George Lowenstein, Ted O'Donoghue, and Matthew Rabin, in their joint contribution to the issue, delineate a variety of legal mechanisms, short of outright proscription, by which the state might seek to discourage irrational behavior.[19] These mechanisms are "asymmetrically paternalistic," in that they create relatively large benefits for those who would otherwise behave irrationally, but impose relatively small costs for rational actors whose behavior does not need to be improved. Camerer and his coauthors are therefore optimistic about the usefulness of law in encouraging citizen rationality. Their optimism is not shared by Ainslie and Monterosso, who express skepticism about law's ability to combat the intertemporal irrationality created by hyperbolic discounting,[20] nor by Posner, who is skeptical about the efficacy of judicial disapproval of greed in changing greedy preferences.[21]

Second, the fact that a legal institution would have been agreed upon by citizens selecting their institutions in an antecedent position of choice might serve to justify that institution. Rationality, on a "contractarian approach," becomes a standard by which to evaluate legal institutions. Social contract views such as that of Rawls[22] famously appeal to hypothetical rational approval under appropriate conditions as an evaluative criterion for political and social arrangements. McClennen follows this tradition when he suggests that the rationality of avoiding substantial losses to "primary goods" justifies legal protection for those goods in the form of constitutional welfare rights.[23] On the other hand, Matthew Adler argues against the view that rational approvability provides a reason to favor any legal institution to which it applies. Imagine that a governmental official is making a choice that

---

[19] Colin Camerer et al., *Regulation for Conservatives: Behavioral Economics and the Case for "Asymmetric Paternalism,"* 151 U. PA. L. REV. 1211 (2003).

[20] Ainslie & Monterosso, *supra* note 4, at 862.

[21] Posner, *supra* note 15, at 1129.

[22] *See generally* JOHN RAWLS, A THEORY OF JUSTICE (1971).

[23] McClennen, *supra* note 9, at 944-46.

will affect law or legal institutions in some way. Imagine, too, that one of the official's options is rationally approvable by each and every citizen. Specifically, this option maximizes each citizen's preferences, given each citizen's probability information at the time of the choice. Does this fact provide the official a moral reason to choose the option? Adler argues that it does not; his argument casts doubt upon the moral importance of "ex ante efficiency" and more generally seeks to undermine any approach that evaluates possible laws, legal doctrines, institutional structures, and so forth by asking which option citizens would rationally choose.[24]

Third, normative accounts of choice can be directly applied to legal actors. One might criticize particular types of governmental choices as irrational and praise others as rational. For example, if the traditional model provides the correct account of rationality, then legal officials who deviate from the model and instead follow the precautionary principle would be behaving irrationally. This would amount to a telling criticism of that principle. Sunstein not only demonstrates the flaws of the precautionary principle as a normative account of choice, but quite directly and understandably infers from that premise the conclusion that policymakers ought not follow it.

Finally, the articles in this issue may help to shed light on the conflicts that arise between the economic criterion of welfare maximization and the traditional noneconomic content of many legal rules, such as rules that create rights. The reason is that enforcing or honoring a legal right will often turn out to conflict with the demands of welfare maximization. In such a case, how can state officials or citizens have moral reason to enforce or honor the right? Recall Mintoff's answer: certain state commitments or plans to recognize legal rights function as "filters" constraining future official or citizen choices, just as certain rational individual plans render rationally ineligible some future choices on the part of that individual.[25] But can the nontraditional accounts of rationality do any better at integrating rights into a rational actor approach to legal rules? Or does incorporating legal rules that articulate rights force us to abandon the rational actor model altogether? Though the articles in this issue do not, individually or collectively, contain answers to deeper jurisprudential difficulties of this sort, our hope is that they will prepare the ground for a broader discussion on such topics.

---

[24] Matthew D. Adler, *The Puzzle of "Ex Ante Efficiency": Does Rational Approvability Have Moral Weight?*, 151 U. PA. L. REV. 1255 (2003).

[25] Mintoff, *supra* note 8, at 908.

*   *   *   *   *   *