

NBER WORKING PAPER SERIES

PARTIAL PRESCRIPTIONS FOR DECISIONS WITH PARTIAL KNOWLEDGE

Charles F. Manski

Working Paper 14396

<http://www.nber.org/papers/w14396>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

October 2008

This research was supported in part by NSF Grant SES-0549544. I have benefitted from the comments of Ken Binmore, Larry Blume, Buz Brock, Dan Hausman, Francesca Molinari, Jörg Stoye, and Alex Tetenov. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2008 by Charles F. Manski. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Partial Prescriptions For Decisions With Partial Knowledge  
Charles F. Manski  
NBER Working Paper No. 14396  
October 2008  
JEL No. D81

**ABSTRACT**

This paper concerns the prescriptive function of decision analysis. I suppose that an agent must choose an action yielding welfare that varies with the state of nature. The agent has a welfare function and beliefs, but he does not know the actual state of nature. It is often argued that such an agent should adhere to consistency axioms which imply that behavior can be represented as maximization of expected utility. However, our agent is not concerned the consistency of his behavior across hypothetical choice sets. He only wants to make a reasonable choice from the choice set that he actually faces. Hence, I reason that prescriptions for decision making should respect actuality. That is, they should promote welfare maximization in the choice problem the agent actually faces. I conclude that any decision rule respecting weak and stochastic dominance should be considered rational. Expected utility maximization respects dominance, but it has no special status from the actualist perspective. Moreover, the basic consistency axiom of transitivity has a clear normative foundation only when actions are ordered by dominance.

Charles F. Manski  
Department of Economics  
Northwestern University  
2001 Sheridan Road  
Evanston, IL 60208  
and NBER  
cfmanski@northwestern.edu

## 1. Introduction

I consider here the prescriptive function of decision analysis, which is to develop constructive rules for decision making. I suppose that an agent—perhaps a firm, an individual, or a planner—must choose an action yielding welfare that depends on the state of nature. The agent has a welfare function and beliefs, which I take as primitives. Formally, the agent knows that the choice set is  $C$ , the feasible states of nature are  $S$ , and the objective is to maximize a function  $w(\cdot, \cdot): C \times S \rightarrow \mathbb{R}^1$  mapping actions  $a \in C$  and states  $s \in S$  into welfare. For example,  $w(\cdot, \cdot)$  may be the profit function of a firm, the utility function of a consumer, or the welfare function of a social planner. The problem is to choose an action without knowing the actual state of nature.

Decision analysis has long distinguished two polar versions of this problem. In choice under uncertainty, the agent only knows  $(C, S, w)$ . In choice under risk, he also knows that the states of nature form a probability space  $(S, \Omega, P)$  and that the actual state of nature will be drawn at random from distribution  $P$ . In both cases, the prevalent prescription has been that the agent should maximize expected utility. Expectations under risk are taken with respect to  $P$ . In the case of uncertainty, the agent is called upon to form a subjective probability space, say  $(S, \Omega_u, P_u)$  and take expectations with respect to  $P_u$ .

Why should an agent maximize expected utility? A standard response has been to cite representation theorems deriving the expected utility criterion from consistency axioms on hypothetical choice behavior, famously Von Neumann and Morgenstern (1944) for risk and Savage (1954, 1972) for uncertainty. However, the N-M and Savage theorems, as well as subsequent work in axiomatic decision theory, do not prescribe how an agent should behave when facing an actual choice problem. Rather, these theorems consider an agent who has formed a complete binary preference ordering over a specified class  $A$  of actions and, thus, who knows how he would behave if he were to face any choice set  $D \subset A$ . The theorems show that if the preference ordering has certain properties, then the agent may be represented as maximizing expected utility. Thus, the theorems of axiomatic decision theory are interpretative rather than prescriptive.

Why then are the N-M and Savage theorems regularly considered to be prescriptive? Decision theorists often assert that an agent *should* form a complete binary preference ordering on the class A of actions and that preferences *should* have the properties assumed in the theorems. If one accepts these assertions, the theorems imply that the agent *should* behave in a manner representable as maximization of expected utility. Thus, the theorems are prescriptive if one considers their consistency axioms to be compelling.

Why should one consider the N-M or Savage axioms to be compelling? No theorem answers this question. Instead, decision theorists call for introspection. Savage (1954, 1972, p. 7) put it this way:

I am about to build up a highly idealized theory of the behavior of a “rational” person with respect to decisions. In doing so I will, of course, have to ask you to agree with me that such and such maxims of behavior are “rational.” In so far as “rational” means logical, there is no live question; and, if I ask your leave there at all, it is only as a matter of form. But our person is going to have to make up his mind in situations in which criteria beyond the ordinary ones of logic will be necessary. So, when certain maxims are presented for your consideration, you must ask yourself whether you try to behave in accordance with them, or, to put it differently, how you would react if you noticed yourself violating them.

Similarly, in lecture notes for a Ph.D. course in decision theory, Kreps (1988, p. 5) counseled an agent contemplating application of the N-M theorem that he must first “Decide that you want to obey the axioms because they seem reasonable guides to behavior.”

I write this paper to report the results of my own introspection on rational behavior when facing decision problems under risk and uncertainty, as well as intermediate problems with a partially known probability distribution on the states of nature. No theorem can prove that my introspection is superior or inferior to that of any axiomatic decision theorist. Why then might someone other than myself find what I have to say to be of interest? I develop two arguments, speaking first to decision making agents and then to

decision theorists.

In Section 2, I speak to an agent who faces the choice problem described in the first paragraph of this paper: the agent wants to maximize  $w$ , but he does not know the actual state of nature. I argue that the standard axiomatic decision theory rationale for expected utility maximization is largely irrelevant to this agent. Standard theory embeds the agent in a world with a class  $A$  of potential actions and asks him to determine how he would behave if he were to face any choice set  $D \subset A$ . The theory begins with the basic axiom that the agent forms a complete binary ordering on  $A$  and then places further consistency axioms on behavior across different choice sets. However, our agent is not concerned with a class of potential actions or with the consistency of his behavior across hypothetical choice sets. He just wants to make a reasonable choice from the choice set that he actually faces. Hence, I reason that prescriptions for decision making should *respect actuality*. That is, they should promote welfare maximization in the choice problem the agent actually faces.

Restricting attention to the actual choice problem, I first consider risk and uncertainty and then extend the discussion to intermediate cases. Under uncertainty, I reason that only one prescription is so compelling as to warrant universal acceptance. That is, the agent should not choose an action that is weakly dominated. Under risk, I propose that a reasonable decision rule should respect both weak and stochastic dominance. Intermediate are choice problems in which the agent knows that the actual state of nature will be drawn from some distribution  $P$ , but only knows that  $P \in \Psi$ , where  $\Psi$  is a given set of probability distributions on  $(S, \Omega)$ . In such settings, I propose that a reasonable decision rule should respect weak dominance and should respect stochastic dominance when the latter property is determinate.

Respect for weak and stochastic dominance are uncontroversial prescriptions for decision making. Both prescriptions respect actuality. However, my position that no other prescription warrants universal acceptance contrasts with the widespread view that an agent should maximize expected utility. The expected utility criterion respects stochastic dominance and it is a minor extension to require respect for weak

dominance. However, I do not see a compelling reason to give this criterion special status. I consider any choice that respects weak and stochastic dominance to be rational. Thus, I offer only partial prescriptions for decisions with partial knowledge.

Section 3 speaks to axiomatic decision theorists who may not accept or appreciate the sharp distinction I draw between hypothetical and actual choice problems. Axiomatic theory is sometimes described as a form of revealed preference analysis, which studies actual choice data rather than statements about preference. However, the enormously rich choice data contemplated in the N-M or Savage axioms are essentially never available in practice. Hence, decision theorists perform thought experiments. Referring to two actions labeled  $f$  and  $g$ , Savage (1954, 1972) wrote (p. 17): “I think it of great importance that preference, and indifference, between  $f$  and  $g$  be determined, at least in principle, by decisions between acts and not by response to introspective questions.”

From my perspective, the critical phrase in this sentence is “at least in principle.” The collection of choices contemplated in the Savage and other axiom systems are hypothetical. They are not observed in practice. This has been pointed out repeatedly over the years, at least as early as Sen (1973). Nevertheless, axiomatic decision theorists persist in referring to their subject as revealed preference analysis.

Section 3 attempts to interpret decision theoretic consistency axioms in terms of actual choice problems, focusing on the axiom that binary choices should be transitive. I observe that the three binary choice problems of a standard transitivity axiom have at least three interpretations—these problems could be mutually exclusive, or the agent could face them simultaneously or sequentially. Considering the first two interpretations, I find that binary choices should be transitive when actions are ordered by dominance but need not be otherwise. I do not formally study sequential choice, but I do discuss the unusual version of sequential choice set forth in the famous “money pump” story, which has often been used to argue for transitivity.

Section 4 makes concluding remarks.

## 2. Choice with an Indeterminate Objective Function

### 2.1. Basic Considerations: Respect for Actuality

The agent under consideration aspires to maximize the objective function  $w(\cdot, r)$ , where  $r$  denotes the actual state of nature. He knows that  $r$  lies in a set  $S$  of feasible states of nature. If there exists no  $a \in C$  that uniformly maximizes  $w(\cdot, s)$  for all  $s \in S$ , his objective function is indeterminate.

Before going further, I should call attention to the fact that research in decision analysis exhibits two starkly different perspectives on objective functions. In axiomatic theory, objective functions are mathematical constructs derived from an agent's preference ordering on the class  $A$  of potential actions. In applied decision analysis, objective functions are psychological constructs that agents use to make choices. I necessarily take the psychological realist position here. If objective functions are constructs inferred from preferences, then prescriptions for decision making are unnecessary as the agent has already determined how he would behave.

The psychological realist perspective is evident in the standard economic theory of the firm. The theory of the firm does not derive a firm's objective function from a preference ordering on potential actions. It assumes that the firm wants to maximize profit.

Throughout this paper I will repeatedly cite a problem of firm decision making that emphasizes the practical nature of my concerns. I will consider a farmer who must choose what crop to plant on a plot of land; the feasible actions may, for example, be to plant corn, soy, or let the land lie fallow. The farmer aspires to maximize profit  $p_{ar}O_{ar} - K_{ar}$ , where  $K_{ar}$  is the cost of planting and harvesting crop  $a$  in state of nature  $r$ ,  $O_{ar}$  is the resulting output, and  $p_{ar}$  is the price obtained per unit of output. Then  $w(c, r) = p_{cr}O_{cr} - K_{cr}$  is the actual value of the objective function when  $c$  is chosen. Suppose, as is realistic, that the farmer must choose a crop with incomplete knowledge of costs, outputs, and prices. Then his objective function

is indeterminate.

Say that the farmer explains his situation and asks for advice on decision making. My thinking on how to respond begins with an idea that a practical decision maker may think so self-evident as not to require mention, but which I will give a formal name:

*Respect for Actuality* (RA): Any prescription for decision making should promote welfare maximization in the choice problem the decision maker actually faces.

In the context of this paper, respect for actuality implies that any prescription for decision making should refer only to the potential values of the agent's objective function, these being  $[w(a, s), a \in C, s \in S]$ , and to available probabilistic information concerning  $S$ . It should not refer to other choice sets or other sets of states of nature. After all, our agent wants to maximize  $w$  in the choice setting he actually faces.<sup>1</sup>

Consider, for example, the farmer planting a crop. The farmer wants to maximize actual profits, whose potential values are  $(p_{as}O_{as} - K_{as}, a \in C, s \in S)$ . A prescription for planting respects actuality if it refers only to these potential profits, not to profits that the farmer might realize if he were to face a different menu of crops or if he were to have different knowledge of costs, outputs, and prices.

With the exception of a brief discussion of sequential choice in Section 3, this paper restricts attention to one-shot choice problems. Readers concerned with sequential choice should not interpret respect for actuality as demanding that agents be myopic. If an agent faces a sequence of choice problems and his welfare depends on the entire sequence, consideration of potential future choices is germane to the choice that the agent makes today and, hence, does respect actuality.

---

<sup>1</sup> Respect for actuality resembles, but is formally distinct from, the *conditionality principle* espoused by Bayesian decision theorists. They argue that statistical inference should be conditioned on observed data alone and should not rest on thought experiments that contemplate how an inferential procedure would perform in hypothetical repeated sampling. See, for example, Berger (1985), Chapter 1.



*Respect for Actuality and Consistency Axioms*

I think it essential to explicitly name *Respect for Actuality* because this idea is deeply at odds with standard axiomatic decision theory, which poses consistency axioms on behavior across hypothetical choice sets. Decision theorists consider adherence to consistency axioms to be a virtue per se. It is revealing to recall how Savage (1954, 1972) viewed the matter. After discussing the positive role of logic in guiding actual human behavior, he wrote (p. 20):

The principal value of logic, however, is in connection with its normative interpretation, that is, as a set of criteria by which to detect, with sufficient trouble, any inconsistencies there may be among our beliefs, and to derive from the beliefs we already hold such new ones as consistency demands. It does not seem appropriate here to attempt an analysis of why and in what contexts we wish to be consistent; it is sufficient to allude to the fact that we often do wish to be so.

Then, addressing his basic postulate P1, which assumes that the agent places a complete binary preference ordering on all potential actions, he wrote

Pursuing the analogy with logic, the main use I would make of P1 and its successors is normative, to police my own decisions for consistency and, where possible, to make complicated decisions depend on simpler ones. Here it is more pertinent than it was in connection with logic that something be said or why and when consistency is a desideratum, though I cannot say much.

He went on to describe a hypothetical conversation with a critic of his theory but, in truth, his brief attempt to motivate his strong concern with consistency does not say much. Considering transitivity, Savage says this to his hypothetical critic (p. 21):

“when it is explicitly brought to my attention that I have shown a preference for  $f$  as compared with  $g$ , for  $g$  as compared with  $h$ , and for  $h$  as compared with  $f$ , I feel uncomfortable in much the same way that I do when it is brought to my attention that some of my beliefs are logically contradictory. Whenever I examine such a triple of preferences on my own part, I find that it is not at all difficult

to reverse one of them. In fact, I find that on contemplating the three alleged preferences side by side that at least one among them is not a preference at all, at any rate not any more.”

This passage appears to describe a psychological process in which Savage uses transitivity as a cognitive tool to learn his own preferences. Sugden (1990) alludes to such a process when he writes (p. 762): “One of the main ways in which we come to know our own preferences is by noting how we in fact choose, or by constructing hypothetical choice problems for ourselves and monitoring our responses.”

Savage also seems to have viewed his consistency axioms as a cognitive tool to learn one’s beliefs. A decision maker adhering to the Savage axioms is representable as placing a subjective probability distribution on the states of nature. Binmore (2008, Section 7.5) interprets Savage as having in mind a “massaging process,” in which a decision maker modifies his hypothetical decisions until he feels comfortable that the implied subjective distribution adequately expresses his beliefs. The idea appears to be that a person holds coherent probabilistic beliefs internally but is psychologically unable to express them directly. Contemplating hypothetical choice problems helps the person discover his internal beliefs.

It is difficult to reconcile this psychological argument with the formal structure of axiomatic decision theory. The theory contemplates a being who arrives with a complete preference ordering, not a cognitively challenged creature who uses thought experiments with hypothetical choice problems to learn about itself. Thus, efforts to motivate adherence to consistency axioms as tools for cognition lie entirely outside of axiomatic theory.

I take no position on the usefulness for cognition of contemplating hypothetical choice problems. In particular, I am agnostic on the usefulness of decision-theoretic consistency axioms as a cognitive tool. To shed light on this requires serious empirical research on cognition; casual introspection does not suffice.

In any case, this paper concerns an agent who knows that he wants to maximize the objective function  $w(\cdot, r)$  on  $C$ . Our agent does not know the state of nature  $r$  that will be realized within the feasible set  $S$ . His problem is partial knowledge of the external world, not incomplete understanding of himself.

### *Respect for Strict Actuality*

I have written that a prescription for decision making that respects actuality should refer only to the potential values of the agent's objective function, namely  $[w(a, s), a \in C, s \in S]$ . Observe that this is an unordered set of welfare values. Thus, I have not required respect for the actual bundling of welfare values into action-specific vectors  $[w(a, s), s \in S], a \in C$ . To recognize this distinction, I will henceforth say that a prescription respects *strict* actuality if it refers only to the welfare vectors  $[w(a, s), s \in S], a \in C$ . The prescriptions to be offered in Section 2.2 respect strict actuality.

To illustrate the difference between respect for actuality and strict actuality, consider the minimax-regret criterion. An agent using this decision criterion chooses an action that solves the problem

$$\min_{c \in C} \max_{s \in S} \{[\max_{a \in C} w(a, s)] - w(c, s)\}.$$

Here  $\max_{a \in C} w(a, s)$  is the maximum welfare attainable in state  $s$  and  $\max_{a \in C} w(a, s) - w(c, s)$  is the regret of action  $c$  in this state. Construction of regret requires unbundling the action-specific welfare vectors to form the state-specific ideal points  $\max_{a \in C} w(a, s), s \in S$ . Thus, the minimax-regret criterion respects actuality, but it does not respect strict actuality.<sup>2</sup>

I do not view respect for strict actuality as a sine qua non for a reasonable decision criterion. Hence, I will not use this concept to argue against the minimax-regret criterion. To the contrary, in Section 2.3 I will cite minimax regret as one of various decision criteria that an agent might reasonably use to choose among the actions that remain after applying the prescriptions of Section 2.2.<sup>3</sup>

---

<sup>2</sup> I am grateful to Larry Blume for raising this issue with me.

<sup>3</sup> Chernoff (1954) argued against minimax regret because the criterion violates the consistency axiom called independence of irrelevant alternatives (IIA). The IIA axiom holds that if an agent is not willing to choose a given action from a hypothetical choice set, then he should not be willing to choose it from any

## 2.2. Respect for Weak and Stochastic Dominance

The remainder of this paper explores the nature of a decision theory grounded in respect for actuality. I first ask what prescriptions for decision making respect actuality and are so compelling as to warrant universal acceptance as tenets of rational behavior. I consider choice under uncertainty and risk to begin, and then broaden the discussion to settings with partial knowledge of a probability distribution on the set of states of nature.

### *Choice under Uncertainty and Risk*

In problems of choice under uncertainty, the only prescription that I think warrants universal acceptance is respect for weak dominance. By definition, action  $c \in C$  is weakly dominated if there exists a  $d \in C$  such that  $w(d, s) \geq w(c, s)$  for all  $s \in S$  and  $w(d, s) > w(c, s)$  for some  $s \in S$ . The prescription is

*Respect for Weak Dominance (RWD):* The agent should not choose a weakly dominated action.

Respect for weak dominance respects strict actuality. The prescription is uniquely compelling because weak dominance defines the circumstances in which an agent wanting to maximize  $w$  knows that choice of one action improves on choice of another. If  $c$  is undominated, no action surely improves on  $c$ . I view choice

---

larger hypothetical choice set; thus, for any  $c \in D \subset E$ , an agent who would not choose  $c$  from  $D$  should not choose  $c$  from  $E$ . Chernoff wrote (p. 426):

A third objection which the author considers very serious is the following. In some examples, the min max regret criterion may select a strategy  $d_3$  among the available strategies  $d_1, d_2, d_3$ , and  $d_4$ . On the other hand, if for some reason  $d_4$  is made unavailable, the min max regret criterion will select  $d_2$  among  $d_1, d_2$ , and  $d_3$ . The author feels that for a reasonable criterion the presence of an undesirable strategy  $d_4$  should not have an influence on the choice among the remaining strategies.

This passage is the totality of Chernoff's argument. He introspected and concluded that any reasonable decision criterion should adhere to IIA, without explaining why he felt this way. He did not argue that minimax-regret decisions have adverse welfare consequences.

of any such action as rational.

In problems of choice under risk, the first question is whether to exploit the known probability distribution on the states of nature. An agent might reason that he wants to achieve high welfare ex post and, hence, should base decisions only on definite knowledge of ex post welfare. If so, he would ignore all probabilistic information. This reasoning is coherent and I would resist calling it irrational. Indeed, a person making a maximin decision behaves in this manner, and I would not call such a person irrational.

Nevertheless, I will presume that our agent uses available probabilistic information, reasoning that he should seek to maximize the ex ante probability of achieving high welfare ex post. If so, the agent should choose an action that respects stochastic dominance. By definition, action  $c \in C$  is stochastically dominated if there exists a  $d \in C$  such that  $P[w(d, s) \leq x] \leq P[w(c, s) \leq x]$  for all  $x \in \mathbb{R}^1$  and  $P[w(d, s) \leq x] < P[w(c, s) \leq x]$  for some  $x \in \mathbb{R}^1$ . The prescription is

*Respect for Stochastic Dominance (RSD)*: The agent should not choose an action that is stochastically dominated.

Respect for stochastic dominance respects strict actuality. Presuming probabilistic information is to be used at all, the prescription is uncontroversial. Stochastic dominance defines the circumstances in which choice of one action probabilistically improves on choice of another. If  $c$  is not stochastically dominated, no other action uniformly improves on  $c$  probabilistically. I view choice of any such action as rational.

When welfare takes only two values, RSD implies a complete binary ordering of actions. Without loss of generality, let the two feasible values of  $w$  be zero and one. Then action  $c$  is stochastically dominated by  $d$  if and only if  $P[w(d, s) = 1] > P[w(c, s) = 1]$ . Hence, the agent should choose an action that maximizes  $P[w(\cdot, s) = 1]$  and should be indifferent among actions that attain the maximum.

When welfare takes multiple values, RSD generally implies a partial rather than complete ordering of actions. Diverse decision criteria may be used to choose among the actions that are not dominated. Proponents of expected utility maximization argue that only this criterion should be used. The expected utility criterion respects strict actuality, but the decision theoretic arguments made for its supremacy concern consistency of behavior across hypothetical choice problems. These arguments have no standing here. From the perspective of this paper, expected utility maximization is an acceptable decision criterion, but it has no special status. For example, an agent might just as well maximize quantile utility, a criterion studied in Manski (1988).

It is well known (e. g., Hanoch and Levy, 1969) that distribution  $F$  stochastically dominates distribution  $G$  if and only if  $\int u(x)dF(x) \geq \int u(x)dG(x)$  for every increasing function  $u(\cdot)$  and  $\int u(x)dF(x) > \int u(x)dG(x)$  for some increasing  $u(\cdot)$ . This result might be thought to imply a conceptual connection between stochastic dominance and expected utility maximization. However, the result is only one of various representation theorems for stochastic dominance. Let  $Q_\alpha(F)$  and  $Q_\alpha(G)$  denote the  $\alpha$ -quantiles of  $F$  and  $G$ . Another representation theorem (e. g. Levy and Kroll, 1978) is that  $F$  stochastically dominates  $G$  if and only if  $Q_\alpha(F) \geq Q_\alpha(G)$  for every  $\alpha \in (0, 1)$  and  $Q_\alpha(F) > Q_\alpha(G)$  for some  $\alpha \in (0, 1)$ . Thus, stochastic dominance is no more closely tied to expected utility maximization than it is to quantile utility maximization.

It is worth noting that the RWD and RSD prescriptions are invariant with respect to ordinal transformations of the objective function  $w(\cdot)$ . From the perspective of this paper, invariance to ordinal transformations is not a property of inherent value—our agent wants to maximize  $w(\cdot)$ , not some transformation of  $w(\cdot)$ . Nevertheless, the property deserves mention given the important role that ordinality has played in decision theory.

*Choice with Partial Probabilistic Knowledge*

Discussions of choice under risk have traditionally supposed that distribution  $P$  is completely known to the agent. Decision theoretic expositions refer to lotteries, coin tosses, balls drawn at random from urns of known composition, and other explicit randomizing devices. Economic applications refer to cross-sectional frequency distributions, such as the distribution of demand among heterogeneous consumers, or to stochastic processes, such as the distribution of unanticipated shocks in a macroeconomic model. In such applications, an agent who knows  $P$  is said to have rational expectations.

It is straightforward to extend the RSD prescription for choice under risk to settings in which  $P$  is an objective probability distribution whose identity is partially known to the agent. Suppose the agent knows that  $P \in \Psi$ , where  $\Psi$  is a given set of distributions on  $(S, \Omega)$ . Then RSD extends as follows.

*Respect for Stochastic Dominance (RSD):* The agent should not choose an action that is stochastically dominated under all feasible distributions on the states of nature.

The extended RSD prescription generically yields only a partial ordering of actions, even when welfare takes just two values. If  $c$  is stochastically dominated by  $d$  for all  $P \in \Psi$ , the agent knows that  $d$  probabilistically improves on  $c$ . If  $c$  is stochastically dominated for some feasible  $P$  but not for others, he cannot conclude that another action probabilistically improves on  $c$ . I view choice of any such action as rational.

A classic decision theoretic exposition of a choice problem with partial knowledge of  $P$  was given by Ellsberg (1961), who supposed that balls may be drawn at random from one urn of known composition or another urn of unknown composition. In the simplest of Ellsberg's thought experiments, balls have two colors, say red (R) and black (B), and there are 100 balls in each urn. The set of feasible states of nature is  $S = \{R, B\} \times \{R, B\}$ , where the first and second bracketed expressions denote the possible colors of balls

drawn from the first and second urn, respectively. The objective probability distribution on  $S$  is  $P = P_1 \times P_2$ , where  $P_1$  and  $P_2$  are the distributions of colors in each urn. The agent is told  $P_1$ , where  $0 < P_1(R) < 1$ , but he is not told  $P_2$ . Hence,  $\Psi = P_1 \times \Psi_2$ , where  $\Psi_2$  are the 101 feasible distributions placing probabilities  $P_2(R) \in \{0, 1/100, 2/100, \dots, 99/100, 1\}$  on drawing a red ball from the second urn. The agent chooses an urn from which to draw a ball. He receives a positive payoff if one color, say red, is drawn and zero payoff otherwise. In this setting, neither action stochastically dominates the other for all  $P \in \Psi$ ; it may be that  $P_1(R) > P_2(R)$  or  $P_1(R) < P_2(R)$ . Hence, the RSD prescription places no restriction on behavior. I view choice of either urn as rational.

Although the RSD prescription has no bite in the Ellsberg setting, it does in other problems with partial probabilistic knowledge. To illustrate, I return to the farmer's planting decision.

*Illustration:* Consider the farmer with potential profits  $(p_{as}O_{as} - K_{as}, a \in C, s \in S)$ . At the time of the planting decision, the farmer may know costs but not outputs and prices. Outputs, which are determined primarily by weather conditions, may have a known objective probability distribution. However, the farmer may feel unable to place a probability distribution on prices, which are determined by poorly understood world market conditions. Hence, the farmer may place a partial probability distribution on profits.

As earlier, let the three feasible actions be to plant corn (action c), plant soy (action d), and let the land lie fallow (action e). Let action e be known to have zero cost, output, and price. Let actions c and d have known costs  $K_{cr} = K_{dr} = 1$ , and known uniform output distributions  $P_{Oc} = U[0, 1]$  and  $P_{Od} = U[0.5, 1]$ . Regarding prices, the farmer only knows that they lie in the bounded range  $(p_c, p_d) \in [2, 4] \times [2, 3]$ .

In this setting, the farmer has enough information to conclude that planting soy stochastically dominates letting the land lie fallow. The profit distribution for planting soy is one of the uniform distributions  $U[0.5p_d - 1, p_d - 1]$ ,  $p_d \in [2, 3]$ . All of these distributions stochastically dominate the profit distribution for letting the land lie fallow, which places probability one at zero profit. Hence, by RSD, the



farmer should not let the land lie fallow.

Planting corn and soy are undominated actions. The profit distribution for planting corn is one of the uniform distributions  $U[-1, p_c - 1]$ ,  $p_c \in [2, 4]$ . In the worst case,  $U[-1, 1]$ , planting corn is surely stochastically dominated by planting soy. However, in the best case,  $U[-1, 3]$ , planting corn is not stochastically dominated by any of the feasible profit distributions for soy.  $\square$

### *Choice with Subjective Probabilistic Knowledge*

It remains to discuss choice with subjective probabilistic knowledge. As with objective functions, research in decision analysis exhibits two perspectives on subjective probability distributions. In the axiomatic theory of Savage (1954), subjective distributions are mathematical constructs derived from an agent's preference ordering on the class of potential actions. In applied decision analysis, subjective distributions are psychological constructs that agents use to make choices. Arguing for the psychological realism of subjective probabilities, Tversky and Kahneman (1974, p. 1130) made plain the difference between the two perspectives, writing:

It should perhaps be noted that, while subjective probabilities can sometimes be inferred from preferences among bets, they are normally not formed in this fashion. A person bets on team A rather than on team B because he believes that team A is more likely to win; he does not infer this belief from his betting preferences. Thus, in reality, subjective probabilities determine preferences among bets and are not derived from them, as in the axiomatic theory of rational decision.

As with objective functions, I take the psychological realist position.<sup>4</sup>

Suppose then that our agent holds the subjective belief that the actual state of nature is drawn at random from a distribution  $P_u$  lying in a set  $\Psi_u$  of such distributions. In traditional Bayesian thinking  $\Psi_u$  is

---

<sup>4</sup> I should acknowledge that I take the psychological realism of subjective probability distributions very seriously, having devoted considerable effort to measuring expectations probabilistically in survey research (Manski, 2004).

a singleton, but a diverse literature considers settings in which  $\Psi_u$  is a set of distributions; see Walley (1991). Suppose that the agent uses his probabilistic beliefs in decision making, reasoning that he should seek to maximize the subjective probability of achieving high welfare. Then, just as in the case of risk, the agent should choose an action that respects stochastic dominance. I consider any such choice to be rational.

### 2.3. Choosing an Undominated Action

Suppose that the agent has eliminated from consideration all actions that are weakly or stochastically dominated. How might he choose among the remaining undominated actions? I write “might” rather than “should” because RWD and RSD exhaust the prescriptions that I believe warrant universal acceptance.

Proponents of Bayesian decision theory counsel the agent to maximize expected utility. Rejecting this as a universal prescription, I could say that a decision analyst should simply list the actions that respect weak and stochastic dominance, and leave it at that. However, I think that decision analysis can be most useful to decision makers if, beyond listing the undominated actions, it seeks to describe in neutral terms the properties of a menu of alternative criteria that have well-understood properties.

I envision using a decision criterion only to choose among the subset of actions that remain after dominated actions have been eliminated. Decision theorists often study the use of a criterion to choose among the complete choice set  $C$ . It is then common to find that a criterion may yield multiple potential choices, some of which are dominated. For example, a weakly dominated action maximizes expected utility if the set of states on which strict dominance occurs has probability zero. A weakly dominated action is a maximin choice if the set of states on which strict dominance occurs do not minimize welfare. Such pathological properties are prevented by using a criterion only to choose among undominated actions.

An example of what I have in mind is my own study (Manski, 1988) of the quantile utility model, where the agent maximizes a specified quantile of the distribution of utility, and the utility mass model,

where he maximizes the probability that utility exceeds a specified threshold. Comparison of these decision criteria with expected utility maximization is revealing. Whereas the latter criterion is invariant only to cardinal transformations of the objective function, the former are invariant to ordinal transformations. Whereas the latter conveys risk preferences through the shape of the utility function, the former do so through the specified quantile or utility threshold, with higher values conveying more risk preference. Although I find the quantile-utility and utility-mass criteria to be intellectually interesting and worthy of consideration in practice, I was careful not to claim them to be superior to other criteria. Instead, I put it this way (p. 82): “I do not think of the ordinal utility models studied here as dominating the expected utility model. Expected utility, quantile utility, and utility mass should, I feel, be thought of as three tightly specified models. Each illuminates a possible mode of rational behavior under uncertainty.”

I similarly find it intellectually interesting and potentially useful for practice to study criteria that use no probabilistic information, such as the maximin and minimax-regret criteria, and hybrids that use partial probabilistic information, such as the maximin expected utility (aka  $\Gamma$ -maximin) and minimax expected regret (aka  $\Gamma$ -minimax regret) criteria. The goal should be to provide agents facing actual choice problems with a menu of criteria for their consideration. For example, Manski (2006) studied a reasonably realistic social planning problem with partial knowledge. I first showed how to identify weakly dominated policies and then characterized Bayesian, maximin, and minimax-regret policies.

Offering an agent a menu of decision criteria asks him to face a second-order choice problem. Whereas the first-order problem is to choose an action from the set  $C$ , now the agent needs to choose a criterion to apply to this first-order problem. It is tempting to offer guidance to the agent on how he might approach the second-order problem. I will not do so here. The first-order choice problem specified at the beginning of this paper states what the agent wants to accomplish and what he knows about the actual state of nature. However, it does not provide a foundation to advise an agent how to choose among undominated actions.

Consider, for example, the farmer who knows that the profit distribution for planting corn is one of the distributions  $U[-1, p_c - 1]$ ,  $p_c \in [2, 4]$ , and that for planting soy is one of the distributions  $U[0.5p_d - 1, p_d - 1]$ ,  $p_d \in [2, 3]$ . As I see it, the farmer could rationally plant either crop. A decision analyst might usefully describe to the farmer the potential profit implications of using different decision criteria to make the choice. However, I do not think that the decision analyst should prescribe a particular criterion.

### 3. Actualist Interpretations of Transitivity

I have sharply differentiated an agent's actual choice problem from the hypothetical problems of axiomatic decision theory. To make sense of axiomatic theory from the perspective of this paper, one needs to interpret decision theoretic thought experiments in terms of actual choice problems. This is not a straightforward matter. Attempting to deconstruct the abstract mathematical statements of standard consistency axioms, I have found that they do not specify essential aspects of the choice problems under consideration. Hence, the axioms have multiple interpretations.

To explain, I consider the venerable axiom that binary choices should be transitive. The analysis in this section shows that transitivity is not a universally compelling prescription when framed in terms of an actual choice problem of the type considered in this paper. A rational decision maker should exhibit transitivity when actions are ordered by dominance. However, transitivity is optional otherwise.

Here is a standard choice-function statement of transitivity: If an agent would uniquely choose  $c$  from set  $(c, d)$  and would uniquely choose  $d$  from  $(d, e)$ , then he should uniquely choose  $c$  from  $(c, e)$ . Although the transitivity axiom seems straightforward, the thought experiment contemplating three binary choice sets  $(c, d)$ ,  $(d, e)$ , and  $(c, e)$  has multiple interpretations in terms of an actual choice problem. These include

*Simultaneous Binary Choice Problems:* The agent simultaneously faces all three choice sets and chooses an action from each one.

*Mutually Exclusive Binary Choice Problems:* The agent will face one of the three choice sets. He must commit to choose an action from each set, should it be the one he will face.

*Sequential Binary Choice Problems:* The agent sequentially faces the three choice sets, the sequence being (c, d), (d, e), (c, e). This interpretation encompasses multiple informational cases. When the agent faces set (c, d), he may or may not know that he will later face (d, e) and (c, e). At the time of each later choice, he may or may not have new knowledge of the state of nature.

The first two interpretations are formally representable in terms of the choice problem studied in this paper. Moreover, one informational case of sequential choice is equivalent to simultaneous choice. In what follows, I flesh out each interpretation in turn. For simplicity, I assume below that the agent has no probabilistic information concerning the set of states of nature.

### 3.1. Simultaneous Binary Choice Problems

Simultaneous choice from multiple choice sets is equivalent to a single choice problem where actions have multiple components. To formalize the problem, let the choice set be  $C = (c, d) \times (d, e) \times (c, e)$ . Thus,  $C$  contains eight feasible actions, each being a three-element vector of sub-actions.

With  $C$  defined this way, the analysis of Section 2 holds as given. The set of states of nature is  $S$ , the objective function is  $w(\cdot, \cdot): C \times S \rightarrow \mathbb{R}^1$ , and the agent wants to maximize  $w$ . The agent should choose an action that respects weak dominance. Any such action is rational.

Adherence to transitivity means that the agent does not choose action  $(c, d, e)$  or  $(d, e, c)$ . If the welfare function is non-separable in the sub-actions, there is no reason to expect transitivity. I will therefore suppose that  $w$  has the additively separable form<sup>5</sup>

$$w[(i, j, k), s] = g(i, s) + g(j, s) + g(k, s),$$

where  $(i, j, k) \in C$  and  $g(\cdot, \cdot): (c, d, e) \times S \rightarrow \mathbb{R}^1$ .

For example, consider the farmer's planting decision. Suppose that the farmer has three identical fields. In one he can plant either corn  $(c)$  or soy  $(d)$ . In the second he can plant soy or let the land lie fallow  $(e)$ . In the third he can plant corn or let the land lie fallow. The farmer wants to maximize total profit. Then  $g$  is the profit function per field, and  $g(i, s) + g(j, s) + g(k, s)$  is the total profit associated with action  $(i, j, k)$  in state of nature  $s$ .

Suppose first that  $g(c, s) \geq g(d, s) \geq g(e, s)$  for all  $s \in S$ ,  $g(c, s) > g(d, s)$  for some  $s \in S$ , and  $g(d, s) > g(e, s)$  for some  $s \in S$ . Then, for all  $(i, j, k) \neq (c, d, c)$ , we have  $w[(c, d, c), s] \geq w[(i, j, k), s]$  for all  $s \in S$  and  $w[(c, d, c), s] > w[(i, j, k), s]$  for some  $s \in S$ . Hence, action  $(c, d, c)$  weakly dominates all other actions. The agent should therefore choose  $(c, d, c)$ , which is transitive.

The agent should similarly choose action  $(d, e, e)$  if  $g(e, s) \geq g(d, s) \geq g(c, s)$  for all  $s \in S$ ,  $g(e, s) > g(d, s)$  for some  $s \in S$ , and  $g(d, s) > g(c, s)$  for some  $s \in S$ . This choice also is transitive.

The above specifications for  $g$  are special—the agent knows enough to determine a welfare-maximizing action. It appears difficult to obtain a succinct characterization of weak dominance for other specifications of  $g$ . However, it is easy to determine dominance computationally when the set  $S$  is finite. Consideration of a simple case with two feasible states of nature suffices to show that intransitive actions

---

<sup>5</sup> Additivity simplifies some computations but is not essential. The discussion below applies if  $w$  has the non-additive form  $w[(i, j, k), s] = F[g(i, s), g(j, s), g(k, s)]$ , where  $(i, j, k) \in C$ ,  $g(\cdot, \cdot): (c, d, e) \times S \rightarrow \mathbb{R}^1$ , and  $F(\cdot, \cdot, \cdot): \mathbb{R}^3 \rightarrow \mathbb{R}^1$  is strictly increasing in each argument.

may be undominated.

### *A Farming Illustration*

Consider the farmer's planting decision with three identical fields, as described above. Let there be two feasible states of nature, with  $s = 1$  being advantageous for planting corn and  $s = 2$  for planting soy. In particular, let the field-specific profit function be:

$g(c, 1) = 2$	$g(d, 1) = -1$	$g(e, 1) = 0$
$g(c, 2) = -1$	$g(d, 2) = 2$	$g(e, 2) = 0$

Then the total profit function is

$w[(c, d, c), 1]$	$w[(c, d, e), 1]$	$w[(c, e, c), 1]$	$w[(c, e, e), 1]$	$w[(d, d, c), 1]$	$w[(d, d, e), 1]$	$w[(d, e, c), 1]$	$w[(d, e, e), 1]$
3	1	4	2	0	-2	1	-1
$w[(c, d, c), 2]$	$w[(c, d, e), 2]$	$w[(c, e, c), 2]$	$w[(c, e, e), 2]$	$w[(d, d, c), 2]$	$w[(d, d, e), 2]$	$w[(d, e, c), 2]$	$w[(d, e, e), 2]$
0	1	-2	-1	3	4	1	2

Inspection of the total profit function shows that action  $(c, e, e)$  is dominated by  $(c, d, c)$  and that  $(d, e, e)$  is dominated by  $(d, d, c)$ . The remaining six actions are undominated. Hence, from the perspective of this paper, the farmer can rationally make any planting decision among these six.

I have previously cited the maximin and minimax-regret criteria as possible approaches to choice among undominated actions. In this illustration, use of either criterion implies that the farmer chooses one of the intransitive actions  $(c, d, e)$  or  $(d, e, c)$ . The maximin result is easy to see—the worst-case profit for the intransitive actions is 1, while the worst-case profits for all other actions are non-positive. Computation

of regret shows that the maximum regret of the intransitive actions is 3 and that all other actions have larger maximum regret.<sup>6</sup>

Choice of (c, d, e) or (d, e, c) is not pathological in this illustration. These actions are appealing because they diversify the farmer, planting corn in one field, soy in another, and letting the third field lie fallow. Indeed, diversification provides a broad rationale for choice of an intransitive action in a simultaneous choice problem.<sup>7</sup>

### 3.2. Mutually Exclusive Binary Choice Problems

Mutually exclusive choice with commitment, like simultaneous choice, is equivalent to a single choice problem where actions have three components. Again let the choice set be  $C = (c, d) \times (d, e) \times (c, e)$ . Now choice of action (i, j, k) means that the agent commits to choose i if he should face set (c, d), j if (d, e), and k if (c, e).

Mutually exclusive choice problems may not be as easy to envisage as simultaneous choice problems, which are ubiquitous. However, they arise routinely when agents encode contingent decision rules as algorithms to be executed by automated systems such as industrial robots or financial trading procedures. Such algorithms commit the automated system to behave in specified ways should it face alternative choice problems.

---

<sup>6</sup> In general, maximin and minimax-regret are distinct criteria yielding different choices. However, these criteria are algebraically equivalent when the same maximum welfare is achievable in each state of nature. This is so in the present illustration, where maximum profit in both states of nature equals 4.

<sup>7</sup> Diversification can also rationalize seemingly incoherent decisions, where an agent makes different choices in identical settings. In Manski (2007, Chapter 11; 2008) and elsewhere, I have studied a class of simultaneous choice problems in which a planner with partial knowledge of treatment response chooses a treatment for each observationally identical member of a large population. When there are two treatments and the optimal treatment is not determinate, the minimax-regret treatment rule always diversifies, randomly assigning one treatment to some persons and the other treatment to the rest of the population.



With  $C$  defined as above, we can construct a choice problem of the type studied in Section 2. To do so, we need to augment the original definition of a state of nature by adding a trinary indicator variable  $\ell$  such that  $\ell = 1$  if the agent will actually face set  $(c, d)$ ,  $\ell = 2$  if he will face  $(d, e)$ , and  $\ell = 3$  if  $(c, e)$ . Then the set of feasible states of nature is  $S^* \equiv S \times (1, 2, 3)$ , where  $S$  is the original set of Section 2. Now let the objective function  $w(\cdot, \cdot): C \times S^* \rightarrow \mathbb{R}^1$  be

$$w[(i, j, k), (s, \ell)] = 1[\ell = 1] \cdot g(i, s) + 1[\ell = 2] \cdot g(j, s) + 1[\ell = 3] \cdot g(k, s),$$

where  $(i, j, k) \in C$  and  $g(\cdot, \cdot): (c, d, e) \times S \rightarrow \mathbb{R}^1$ . Thus, welfare is  $g(i, s)$  if  $\ell = 1$ ,  $g(j, s)$  if  $\ell = 2$ , and  $g(k, s)$  if  $\ell = 3$ . Observe that I do not let  $\ell$  be an argument of  $g(\cdot, \cdot)$ . If it were, welfare would be choice-set specific and there would be no reason to expect transitivity.

As with simultaneous choice, the welfare maximizing choice is determinate if  $g(c, s) \geq g(d, s) \geq g(e, s)$  for all  $s \in S$ ,  $g(c, s) > g(d, s)$  for some  $s \in S$ , and  $g(d, s) > g(e, s)$  for some  $s \in S$ . Then the agent should choose action  $(c, d, c)$ , which weakly dominates all other actions. The agent should similarly choose action  $(d, e, e)$  if  $g(e, s) \geq g(d, s) \geq g(c, s)$  for all  $s \in S$ ,  $g(e, s) > g(d, s)$  for some  $s \in S$ , and  $g(d, s) > g(c, s)$  for some  $s \in S$ . These choices are transitive.

As with simultaneous choice problems, intransitive actions may be undominated when the welfare maximizing choice is indeterminate. Consider again the planting decision of Section 3.1, modified so that the farmer will be permitted to plant in only one of the three fields, determined by the realization of  $\ell$ . With  $g$  as specified earlier, the total profit function is

$w[(c, d, c), (1, 1)]$	$w[(c, d, e), (1, 1)]$	$w[(c, e, c), (1, 1)]$	$w[(c, e, e), (1, 1)]$	$w[(d, d, c), (1, 1)]$	$w[(d, d, e), (1, 1)]$	$w[(d, e, c), (1, 1)]$	$w[(d, e, e), (1, 1)]$
2	2	2	2	-1	-1	-1	-1
$w[(c, d, c), (2, 1)]$	$w[(c, d, e), (2, 1)]$	$w[(c, e, c), (2, 1)]$	$w[(c, e, e), (2, 1)]$	$w[(d, d, c), (2, 1)]$	$w[(d, d, e), (2, 1)]$	$w[(d, e, c), (2, 1)]$	$w[(d, e, e), (2, 1)]$
-1	-1	-1	-1	2	2	2	2
$w[(c, d, c), (1, 2)]$	$w[(c, d, e), (1, 2)]$	$w[(c, e, c), (1, 2)]$	$w[(c, e, e), (1, 2)]$	$w[(d, d, c), (1, 2)]$	$w[(d, d, e), (1, 2)]$	$w[(d, e, c), (1, 2)]$	$w[(d, e, e), (1, 2)]$
-1	-1	0	0	-1	-1	0	0
$w[(c, d, c), (2, 2)]$	$w[(c, d, e), (2, 2)]$	$w[(c, e, c), (2, 2)]$	$w[(c, e, e), (2, 2)]$	$w[(d, d, c), (2, 2)]$	$w[(d, d, e), (2, 2)]$	$w[(d, e, c), (2, 2)]$	$w[(d, e, e), (2, 2)]$
2	2	0	0	2	2	0	0
$w[(c, d, c), (1, 3)]$	$w[(c, d, e), (1, 3)]$	$w[(c, e, c), (1, 3)]$	$w[(c, e, e), (1, 3)]$	$w[(d, d, c), (1, 3)]$	$w[(d, d, e), (1, 3)]$	$w[(d, e, c), (1, 3)]$	$w[(d, e, e), (1, 3)]$
2	0	2	0	2	0	2	0
$w[(c, d, c), (2, 3)]$	$w[(c, d, e), (2, 3)]$	$w[(c, e, c), (2, 3)]$	$w[(c, e, e), (2, 3)]$	$w[(d, d, c), (2, 3)]$	$w[(d, d, e), (2, 3)]$	$w[(d, e, c), (2, 3)]$	$w[(d, e, e), (2, 3)]$
-1	0	-1	0	-1	0	-1	0

All actions are undominated. Hence, the farmer can rationally make any planting decision. All actions are maximin choices, as all have worst-case profit equal to  $-1$ . All actions are minimax-regret choices, as all have maximum regret equal to 3.

### 3.3. Sequential Choice Problems

Sequential choice is equivalent to simultaneous choice if the agent knows *ex ante* that he will face all three choice problems and if he receives no new information during the choice process. For example, the farmer with three fields may have a labor or equipment constraint that enables him to plant only one field at a time. Then he may physically plant the three fields sequentially but view planting as a simultaneous multi-field decision problem.

Sequential choice differs from simultaneous choice in other informational settings. When the agent faces later stages of the choice sequence, he may obtain new knowledge that enables him to update his beliefs about the feasible states of nature. When he faces the first choice set  $(c, d)$ , he may not know that he will

later face (d, e) and (c, e). Moreover, the term “not know” has two interpretations. First, the agent may be aware from the outset that he may potentially face later choice problems, but not know whether he will in fact do so. Second, he may believe that he will face no future choice problems and then be surprised when they appear.

The first interpretation of “not know” can be formalized by augmenting the original definition of a state of nature, as we did to specify mutually exclusive choice problems. Here one would add an indicator variable to denote the future choice problems that may possibly arise. The second interpretation requires modification of the concept of a state of nature to permits the agent to find himself in a situation that he earlier thought infeasible. I will not speculate on how to go about this, as decision theorists have long struggled with the question. See, for example, Blume, Brandenburger, and Dekel (1991).

I will not attempt to formalize the numerous possible interpretations of sequential choice here. I will only give an example to illustrate the need for care.

The famous “money pump” argument warning against intransitivity contemplates a sequential choice problem. Binmore (2008) describes the money pump argument this way, referring to a hypothetical decision maker Pandora (p. 11):

Why should we expect Pandora to reveal transitive preferences? The money pump argument says that if she doesn't, other people will be able to make money out of her. Suppose that Pandora's reveals the intransitive preferences  $\text{apple} \prec \text{orange} \prec \text{fig} \prec \text{apple}$  when making pairwise comparisons. A trader now gives her an apple. He then offers to exchange the apple for an orange, provided she pays him a penny. Since her preference for the orange is strict, she agrees. The trader now offers to exchange the orange for a fig, provided she pays him a penny. When she agrees, the trader offers to exchange the fig for an apple, provided she pays him a penny. If Pandora's preferences are stable, this trading cycle can be repeated until Pandora is broke. The inference is that nobody with intransitive preferences will be able to survive in a market context.

Although this familiar argument states that the three trades are offered sequentially, it does not state whether Pandora knows at the time of the first trade that she will later face the second and third ones. Pandora apparently believes that each trade is a one-shot choice problem and does not update this belief when the trader appears a second and third time. When the trading cycle is repeated, Pandora acts as if there was no previous history of trades and there will be no future trades.

It is clear that Pandora's behavior is awfully stupid. However, the money pump argument is just a story about a strangely memoriless and myopic sequential choice process. If a decision maker recognizes that she faces a sequence of trades, she should view herself as facing a simultaneous choice problem rather than isolated binary choice problems. Pandora's poor fate results from her inability to contemplate the future or learn from the past. It is not remotely a general argument for transitivity.

#### 4. Conclusion

I have long felt uneasy about the prescriptive assertion that decision makers *should* maximize expected utility. I could appreciate the mathematical creativity of the N-M, Savage and related representation theorems, but I could not understand why many decision theorists view the consistency axioms of these theorems to be normative. It has long been a mystery to me that the deductive logic of the theorems could be so tight, yet the introspective arguments offered for the axioms could be so loose.

I have occasionally voiced my concerns to close friends, but I have until now not expressed them in writing, other than in brief remarks in Manski (1988). It took this long for me to reach the conclusion that I have something worth writing on the subject. Now I have, the central idea being *respect for actuality*.

When I began work on this paper, I did not anticipate that I would find it difficult to motivate the prescriptive use of such a basic consistency axiom as transitivity. However, this is the way it has worked

out. Supposing that an agent has no probabilistic information on the state of nature, I have found that transitivity has a clear normative foundation only when actions are ordered by weak dominance. It would be valuable to examine other consistency axioms in the same manner.

I stated at the outset and have repeated that this paper concerns the prescriptive function of decision analysis. Nevertheless, it is natural to ask in conclusion whether the work has implications for prediction of behavior. Most economists and some other behavioral scientists maintain the premise in their research that humans aim to be rational in practice. To such researchers, the manner in which prescriptive analysis conceptualizes rationality is relevant to prediction.

An adherent of standard axiomatic theory would predict that real agents maximize expected utility. In contrast, a person who finds respect for actuality to be compelling would predict that agents respect weak and stochastic dominance, but would not take a stand on the empirical prevalence of expected utility maximization. The former person would predict that real agents make transitive binary choices. The latter one would not take a stand on the empirical prevalence of transitivity in settings where actions are undominated. In these and other ways, the present work would thus have implications for prediction of behavior. I cautiously write “would” rather than “does” because this is not the place to evaluate the empirical validity of the premise that humans respect weak and stochastic dominance.

## References

- Berger, J. (1985), *Statistical Decision Theory and Bayesian Analysis*, New York: Springer-Verlag.
- Binmore, K. (2008), *Rational Decisions*, Princeton: Princeton University Press, forthcoming.
- Blume, L., A. Brandenburger, and E. Dekel (1991), "Lexicographic Probabilities and Choice under Uncertainty," *Econometrica*, 59, 61-79.
- Chernoff, H. (1954), "Rational Selection of Decision Functions," *Econometrica*, 22, 422-443.
- Ellsberg, D. (1961), "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics*, 75, 643-669.
- Hanoch, G. and H. Levy (1969), "The Efficiency Analysis of Choices Involving Risk," *Review of Economic Studies*, 36, 335-346.
- Kreps, D. (1988), *Notes on the Theory of Choice*, Boulder: Westview Press.
- Levy, H. and Y. Kroll (1978), "Ordering Uncertain Options with Borrowing and Lending," *Journal of Finance*, 33, 553-574.
- Manski, C. (1988), "Ordinal Utility Models of Decision Making Under Uncertainty," *Theory and Decision*, 25, 79-104.
- Manski, C. (2004), "Measuring Expectations," *Econometrica*, 72, 1329-1376.
- Manski, C. (2006), "Search Profiling with Partial Knowledge of Deterrence." *Economic Journal*, 116, F385-F401.
- Manski, C. (2007), *Identification for Prediction and Decision*, Cambridge, Mass.: Harvard University Press.
- Manski, C. (2008), "Adaptive Partial Policy Innovation: Coping with Ambiguity through Diversification," Department of Economics, Northwestern University.
- Savage, L. (1954), *The Foundations of Statistics*, New York: Wiley.
- Savage, L. (1972), *The Foundations of Statistics*, New York: Dover.
- Sen, A. (1973), "Behaviour and the Concept of Preference," *Economica*, 40, 241-259.
- Sugden, R. (1990), "Rational Choice: A Survey of Contributions from Economics and Philosophy," *The Economic Journal*, 101, 751-785.
- Tversky, A. and D. Kahneman (1974), "Judgement Under Uncertainty: Heuristics and Biases," *Science*, 185, 1124-1131.

von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.

Walley, P. (1991), *Statistical Reasoning with Imprecise Probabilities*, Chapman & Hall: London.