ACTUALIST RATIONALITY

Charles F. Manski
Department of Economics and Institute for Policy Research,
Northwestern University

December 31, 2008

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 an objective 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 argue that any decision rule respecting weak and stochastic dominance is rational from the actualist perspective.

This research was supported in part by NSF Grant SES-0549544. This paper revises and supercedes "Partial Prescriptions for Decisions with Partial Knowledge." I have benefitted from the comments of Ken Binmore, Larry Blume, Buz Brock, Dan Hausman, Francesca Molinari, Jörg Stoye, and Alex Tetenov.

1. Introduction

I consider here the prescriptive function of decision analysis, which is to develop reasonable rules for decision making. I consider a familiar class of problems of decision with partial knowledge. 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 an objective function and beliefs, which I take as primitives. His problem is to choose an action without knowing the actual state of nature.

Formally, the agent faces choice set C and believes that the actual state of nature lies in the set S. The objective function $w(\cdot, \cdot): C \times S \to R^1$ maps actions and states 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 planner. The agent wants to maximize $w(\cdot, r)$, where r is the actual state of nature, but he does not know r. He only knows that $r \in S$.

I will repeatedly cite a problem of firm decision making that emphasizes the practical nature of my concerns. Consider a farmer who must choose what crop to plant on a plot of land. The feasible actions are to plant corn (c), soy (d), or let the land lie fallow (e). Thus, the choice set is $C = \{c, d, e\}$. The farmer wants to maximize profit $w(a, r) = p_{ar}O_{ar} - K_{ar}$, 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. The farmer must choose a crop with partial knowledge of costs, outputs, and prices.

Say that the farmer or another agent explains his decision problem and asks for advice. My thinking on how to respond stems from 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): Prescriptions for decision making should promote welfare maximization in the choice problem the agent actually faces.

In the context of this paper, respect for actuality implies that a prescription for decision making should refer only to the feasible values of the agent's objective function, these being [w(a, s), $a \in C$, $s \in S$]. It should not refer to other choice sets, states of nature, or objective functions. After all, our agent wants to maximize welfare in the choice setting he actually faces. Consider the farmer planting a crop. A prescription for planting respects actuality if it refers only to the feasible profit values ($p_{as}O_{as} - K_{as}$, $a \in C$, $s \in S$). It should not refer to profits that the farmer might hypothetically earn if he were to face a different menu of crops or if he were to have different knowledge of costs, outputs, and prices.

I find it essential to explicitly name *Respect for Actuality* because this idea is deeply at odds with axiomatic decision theory, which has long dominated economic thinking about decision making with partial knowledge. Axiomatic decision theorists consider behavior across hypothetical choice sets and propose that such behavior should exhibit specified forms of consistency, formalized in axioms. Indeed, decision theorists often argue that adherence to certain consistency axioms defines rational behavior.

Consider the famous Von Neumann and Morgenstern (1944) and Savage (1954) representation theorems deriving the expected utility criterion from consistency axioms on hypothetical choice behavior.[1] The N-M and Savage 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 hypothetical choice set $D \subset A$. The theorems show that if the preference ordering has certain consistency properties, then

[1] Axiomatic decision theorists may not accept or appreciate the sharp distinction I draw between hypothetical and actual choice behavior. 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.

the agent may be represented as maximizing expected utility. Thus, the N-M and Savage theorems are interpretative rather than prescriptive.[2]

Nevertheless, the N-M and Savage theorems are regularly considered to be prescriptive. This is so because axiomatic 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 consistency axioms. 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.

---

[2] Whereas I take the objective function w to be a psychological construct that an agent uses to make decisions, the N-M and Savage theorems consider w to be a mathematical construct which can be derived (up to a cardinal transformation) from an agent's preference ordering on the class A of potential actions. I necessarily take the psychological realist position in this paper. 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.

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."

As I see it, consistency axioms are irrelevant to an agent who faces the choice problem described in the opening paragraphs of this paper. This agent wants to maximize a specified objective function, but he does not know the actual state of nature. He is not concerned with the consistency of his behavior across hypothetical choice sets. He just wants to make a reasonable choice from the one choice set that he actually faces. Hence, I argue that prescriptions for decision making should respect actuality.

Section 2 develops principles for actualist rationality. These principles refer only to the choice problem that the agent actually faces, not to hypothetical choice problems. I first consider the basic form of the choice problem, where the agent only knows (C, S, w), and then ones in which he believes that the actual state of nature is drawn at random from some probability distribution on S. In the former setting, 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. In the latter setting, I reason that the agent should also not choose an action that is stochastically dominated.

Respect for weak and stochastic dominance are uncontroversial prescriptions for decision making with partial knowledge. However, my position that no other prescription warrants universal acceptance contrasts with the widespread view among economists 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. Any choice that respects weak and stochastic dominance is rational from the actualist perspective. Thus, I offer only partial prescriptions for decisions with partial knowledge.

My discussion of actualist rationality gives axiomatic decision theory no prescriptive function. Yet Savage and other axiomatic decision theorists have argued that consistency axioms are normative. Seeking to understand the Savage perspective, I have concluded that he had in mind a subtle decision process in which an agent uses axioms as cognitive tools for self-discovery of objectives and beliefs. Section 3 explains, quoting Savage at some length. Section 4 concludes.

2. Actualist Prescriptions for Decisions with Partial Knowledge

Section 2.1 develops actualist prescriptions for decision making by an agent facing the choice problem set out at the beginning of this paper. I first consider the basic setting in which the agent only knows (C, S, w) and then ones in which the agent believes that the actual state of nature is drawn at random from some probability distribution on S. I argue that a reasonable choice should respect weak and stochastic dominance. Section 2.2 considers choice among undominated actions.

2.1. Respect for Weak and Stochastic Dominance

When an agent faces the basic version of the choice problem, 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 actuality. The prescription is uniquely compelling because weak dominance defines the circumstances in which an agent who wants to maximize w knows that choice of one action improves on choice of another. If c is undominated, no action surely improves on c.

Suppose that, in addition to knowing (C, S, w), the agent knows (or believes) that S is a probability space (S, Ω, P) and that the actual state of nature is drawn at random from distribution P. The distinction between "knows" and "believes" has been important in axiomatic decision theory. Von Neumann and Morgenstern took P to be a specified objective probability distribution, while Savage took it to be a subjective distribution which can be inferred from hypothetical choice behavior.[3] However, the distinction is immaterial here.

The first question is whether an agent should exploit knowledge of P in decision making. He 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. Nevertheless, I will presume that the agent uses available probabilistic information, reasoning that he should seek to maximize the ex ante probability of achieving high welfare ex post.

Given this, 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 $P[w(d, s) \le x] \le P[w(c, s) \le x]$ for all $x \in R^1$ and $P[w(d, s) \le x] < P[w(c, s) \le x]$ for some $x \in R^1$. The prescription is

---

[3] Whereas Savage (1954) and subsequent axiomatic decision theorists have viewed subjective distributions as mathematical constructs derived from an agent's preference ordering on the class of potential actions, applied decision analysts have viewed them as 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.

I take the psychological realism of subjective probability distributions very seriously, having devoted considerable effort to measuring expectations probabilistically in survey research (Manski, 2004).

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

Respect for stochastic dominance respects 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.

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 $P[w(\cdot, s) = 1]$.

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 stochastically dominated. Proponents of expected utility maximization argue that only this criterion should be used. However, the axiomatic arguments made for its supremacy concern consistency of behavior across hypothetical choice problems. These arguments have no standing here. From the actualist perspective, expected utility maximization is an acceptable decision criterion but 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 α-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(·). Invariance to ordinal transformations is not a property of inherent value from the actualist perspective—our agent wants to maximize w(·), not some transformation of w(·). Nevertheless, the property deserves mention given the important role that ordinality has played in decision theory.

Choice with Partial Probabilistic Knowledge

Researchers have traditionally supposed that distribution P is completely known to the agent. However, it is straightforward to extend the RSD prescription to settings in which the identity of P is only partially known. Suppose the agent knows that $P \in \Psi$, where Ψ is a given set of distributions on (S, Ω). Then RSD extends as follows.

Respect for Stochastic Dominance (RSD): The agent should not choose an action c if there exists another action d that stochastically dominates c 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 by d for some feasible P but not for others, he cannot conclude that d probabilistically improves on c.

A classic decision theoretic exposition of a choice problem with partial knowledge of P was given by Ellsberg (1961), who supposed that balls are drawn at random from one urn of known composition and 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 Ψ_2 are the 101 feasible distributions placing probabilities $P_2(R) \in \{0, 1/100, 2/100, \ldots, 99/100, 1\}$ 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. Choice of either urn is rational.

Although the RSD prescription has no bite in the Ellsberg setting, it does in other problems with partial probabilistic knowledge. To illustrate, I consider 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 (c), plant soy (d), and let the land lie fallow (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. □

2.2. 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 Γ-maximin) and minimax expected regret (aka Γ-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.[4]

[4] Axiomatic decision theorists have regularly argued against decision criteria other than expected utility maximization on the grounds that they violate one or another consistency axiom. These arguments have no standing from the actualistic perspective.

A famous example is the Chernoff (1954) argument that the minimax regret criterion should not be used because it 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 larger hypothetical choice set; thus, for any $c \in D \subset E$, an agent

Planting Multiple Fields

Here is an extended, hopefully instructive, illustration of what I have in mind. Suppose that a 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. Thus, the farmer's choice set is C = (c, d) × (d, e) × (c, e). This set contains eight feasible actions, each being a three-element vector of sub-actions.

Suppose that the farmer wants to maximize total profit

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

where $(i, j, k) \in C$ and $g(\cdot, \cdot): (c, d, e) \times S \rightarrow R^1$. Here g is the profit function per field in state of nature s. 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 profit function per field 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

---

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.

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

Suppose that the farmer has no probabilistic knowledge of the state of nature. Inspection of the total profit function shows that action (c, e, e) is weakly dominated by (c, d, c) and that (d, e, e) is weakly dominated by (d, d, c). The remaining six actions are undominated. Hence, I would advise the farmer that he can rationally make any planting decision among these six.

An axiomatic decision theorist might argue against choice of (c, d, e) or (d, e, c) on the grounds that these actions are intransitive. 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). This seems straightforward, but the transitivity axiom does not completely specify the thought experiment contemplating choice from the hypothetical choice sets (c, d), (d, e), and (c, e). Three interpretations are

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

Mutually Exclusive 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 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.[5]

The formal structure of the transitivity axiom does not specify which, if any, of these interpretations the axiom is intended to embrace.[6]

In any case, our farmer faces a set of simultaneous choice problems. I have previously cited the

---

[5] 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. 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. Or he may believe that he will face no future choice problems and then be surprised when they appear.

[6] Whereas the formal structure of the axiom is incomplete, accompanying discussions sometimes partially clarify the intended interpretation. The famous "money pump" argument warning against intransitivity contemplates sequential choice problems. Binmore (2009) 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 apple $\prec$ orange $\prec$ fig $\prec$ 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.

This familiar argument states that the three trades are offered sequentially, but 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.

maximin and minimax-regret criteria as possible approaches to choice among undominated actions. 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.[7]

As I see it, choice of (c, d, e) or (d, e, c) is rational. Indeed, these actions are appealing because they diversify the farmer, planting corn in one field, soy in another, and letting the third field lie fallow. Diversification provides a rationale for choice of an intransitive action in a simultaneous choice problem.[8]

3. Consistency Axioms as Cognitive Tools

The discourse on actualist rationality in Sections 1 and 2 gives axiomatic decision theory no prescriptive function. Yet decision theorists have long asserted that axioms are normative. In particular, Savage (1954, 1972) viewed himself as developing a theory of rational behavior with partial knowledge. See the passage quoted in Section 1.

Seeking to understand the Savage perspective, I have concluded that he had in mind a far more subtle decision problem than the one analyzed in this paper. I have assumed that the agent knows his objective

---

[7] 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.

[8] 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.

function and has well-defined beliefs on the states of nature. In contrast, Savage was concerned with decision making when an agent who is not initially conscious of his objective and beliefs strives to learn them through a process of introspection. Thus, Savage viewed his axioms as cognitive tools for self-discovery.

To back up this conclusion, I shall quote Savage at some length. After discussing the positive role of logic in guiding actual human behavior, Savage 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."

The above passage describes 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 (2009, 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.

I find it difficult to reconcile the use of consistency axioms as cognitive tools with the formal structure of axiomatic decision theory. The theory formally 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 formal axiomatic theory.

I will take no position here on the usefulness for cognition of contemplating hypothetical choice problems. In particular, I am agnostic on the usefulness of consistency axioms as cognitive tools. 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 actual state of nature r, but he does know that r lies in the set S. His problem is partial knowledge of the external world, not incomplete understanding of himself.

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*.

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 axiomatic theory would predict that real agents strive to maximize expected utility. In contrast, a person who finds respect for actuality to be compelling would predict that agents aim to respect weak and stochastic dominance, but would not take a stand on the empirical prevalence of expected utility maximization. In this way, the present work would have implications for prediction of behavior.

References

Binmore, K. (2009), *Rational Decisions*, Princeton: Princeton University Press.

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), "Diversified Treatment under Ambiguity," 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.