Reconciling Some Conflicting Evidence on Decision Making under Uncertainty[*]

Edward J. Green

Federal Reserve Bank of Minneapolis

Preliminary Draft

Abstract: Laboratory experiments concerning decision under uncertainty tend to uncover systematic violations of Bayesian rationality. When models that posit Bayesian rationality are compared to non-experimental data, though, they fit the data well. One possible explanation is that an agent's global pattern of choices may not be rationalizable, but that the pattern may satisfy weak conditions sufficent to rationalize the limited range of choices required by any particular decision protocol. Examples of such patterns are constructed here. An agent who adopts a protocol acts rationally, but an experimenter induces irrationality by imposing distinct protocols in various phases of the experiment.

---

[*] This paper reports results of research that was initiated during a visit in June, 1994 to The Institute for Advanced Studies and the Center for the Study of Interactive Rationality at The Hebrew University of Jerusalem. The views expressed in this paper are those of the author, and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Introduction

When people make decisions under uncertainty, are their actions in accord with subjective-expected-utility (SEU) theory? That is, does a person tend to behave as though (*a*) his initial beliefs were encoded by a probability measure on possible events, (*b*) he considers evidence by conditioning this probability measure according to Bayes' rule, and (*c*) he chooses an action that would maximize the expectation of some utility function (which is always the same across the various decision problems that he faces) with respect to this conditional probability measure?

This question has been studied intensively. Researchers who use non-experimental methods tend to conclude that SEU theory accounts well for their data. In sharp contrast, experimental researchers tend to report that behavior inconsistent with SEU theory is pervasive, robust, and easy to elicit from subjects. This systematic disparity between experimental and non-experimental findings is a paradox that needs to be resolved.[1] In this paper I propose one resolution among the many complementary resolutions that can probably account for aspects of this puzzle.

The resolution to be proposed here has to do with how a decision maker often works: by obtaining answers to a sequence of questions, and then eventually choosing an action on the basis of these answers. Such questions typically have answers that are of a yes-or-no variety, or that are numerical quantities, or that otherwise partition the situation into possibilities that are jointly exhaustive and mutually exclusive. Each successive question refines the partition further. It is even common for a decision maker to work according to a protocol that raises these questions in a fixed order. Sometimes an order is determined by practical considerations (as when a detective goes quickly to the scene of the crime, where weather and traffic are likely to obliterate evidence such as footprints), but often a determinate order is followed purely as a matter of professional discipline (as when a doctor obtains a patient's history, if it is possible to do so, at the beginning of treating an illness—even though the nature of the illness is often obvious from signs and symptoms that the patient presents). This determinateness of order has the implication that there is a fixed sequence of increasingly fine partitions of events that uniformly characterize the decision maker's knowledge at successive stages of a family of related decision problems such as criminal investigations or medical cases.

Part of a decision maker's job is to decide when to stop asking questions and to do something. However, even before this point is reached, a decision maker will make judgments regarding what would be the best thing to do on the basis of current evidence. Sherlock Holmes might be confident that he will solve a case conclusively but might nevertheless judge that, if there were an imminent prospect of the criminal committing another heineous crime before further evidence could be unearthed, then the police should try to prevent it by arresting Moriarity. A surgeon might know that, if the hospital had run out of a dye needed to make a patient's blood vessels show clearly in an X-ray, then a particular operation should be performed without waiting to take such an X-ray

---

[1] This paradox has been discussed by the contributors to Hogarth and Reder (1987). In addition to the types of study discussed in that volume, a new kind of non-experimental study has subsequently become available: inference about decision making based on explicit modelling and estimation of data regarding decision making by an expert subject in the course of his actual work. A leading example of such a study is Rust (1987).

because prompt repair would be imperative if the patient were bleeding internally. I will assume that a researcher can observe such "premature" contingent decisions, either because exceptional circumstances such as those that I have just described do arise occasionally, or else because decision makers are able to make reliable reports about what they would do counterfactually. In this respect, I am making a generous assumption about how much data a non-experimental researcher can obtain.

To summarize, in non-experimental settings a decision maker will typically consider a fixed sequence of information partitions, which is dictated either by feasibility restrictions or by the decision maker's protocol. This fixed sequence of partitions is all that a non-experimental researcher can observe. In contrast. an experimental researcher can force a subject to consider information partitions that the researcher—not the subject—selects. In particular, a researcher can force a subject to use incomparable partitions to make decisions about cases that are otherwise similar. For example, a doctor might be required to choose a hypothetical treatment on the basis of information about whether or not the patient has "an elevated temperature" in one case, but on the basis of information about whether or not the patient has "a high fever" in another case. Those two information partitions are intuitively close to one another, but they are not identical. In medicine and in many other fields, authoritative protocols for decision making specify recommended information partitions exactly and discourage decision makers from substituting approximately equivalent partitions. Thus subjects, and especially experts who are well-trained professionals, can be forced in experimental settings to make decisions that they would systematically avoid in their actual work. I am going to argue that comparisons among such forced decisions in experimental settings greatly expand the opportunity to observe violations of the SEU theory. Some implications of this argument will be examined in the concluding section of the paper.

2. A formal description of decision-making data

I will describe the data from experimental and non-experimental studies of decision making in terms that closely resemble a standard version of SEU theory.[2] There is a finite Boolean algebra $\mathcal{E}$, the elements of which are to be interpreted as *events*.[3] I will assume informally that, except for the null event (denoted by $0 \in \mathcal{E}$), these are events to which the decision maker would consider seriously as possibly occurring.[4] Upper-case parameters and variables will denote events in $\mathcal{E}$.

There is also a finite set $\mathcal{A}$ of *acts* that the decision maker can choose. Lower-case parameters and variables will denote acts in $\mathcal{A}$.

---

[2] A significant limitation of this description will be discussed in the conclusion.

[3] Finiteness is assumed here for convenience. It is not a hypothesis of the results of Green and Park (1993) that are adapted to this setting below.

[4] If there are nonzero events that the decision maker would totally discount, then they should constitute a lattice ideal. The Boolean algebra formed by factoring by this ideal will then satisfy my informal assumption. This is the technical reason why I am considering an abstract Boolean algebra of events, rather than a field of sets of "states of nature."

A *decision rule* is a correspondence

$$\delta\colon \mathcal{E} \to \mathcal{A}. \tag{1}$$

That is, δ is a function that assigns a subset of $\mathcal{A}$ to each event. Since $0 \in \mathcal{E}$ is to be interpreted as the impossible event, stipulate that

$$\delta(0) = \mathcal{A}. \tag{2}$$

This stipulation is innocuous, and it will simplify some definitions below.

I will assume that a researcher can observe a subject's entire decision rule in an experiment. All that can be non-experimentally observed, though, is a decision maker's *contingent plan.* A contingent plan consists of a sequence $\Pi = (\Pi_0, \ldots, \Pi_n)$ of partitions[5] of 1, together with a correspondence

$$\pi\colon \bigcup_{m \le n} \Pi_m \to A. \tag{3}$$

Assume that the initial partition Π_0 is trivial, and the partitions Π_m are successive refinements of one another. That is,

$$\Pi_0 = \{1\}, \quad \text{and} \quad \forall m < n \;\; \Pi_{m+1} \text{ is a refinement of } \Pi_m. \tag{4}$$

Given a function π that is defined on the union of a sequence Π of partitions satisfying (4), a sequence Π' satisfying (4) can be constructed from the domain of π. The new Π can differ from the original Π only with respect to the timing of the resolution of uncertainty. (That is, a partition element of some Π_m may be partitioned further at a different time in Π' than in Π, but it must be partitioned in the same way in both sequences.) Since conditional expected utility is insensitive to the time at which uncertainty is resolved, the function π contains all of the information about the sequence Π that will be of interest here. Therefore the term 'contingent plan' will sometimes be applied to π itself, without specifying Π explicitly.

Finally, say that a decision rule δ and a sequence Π of partitions satisfying (4) *induce* the contingent plan π that is obtained by restricting δ to the set of events $\bigcup_{m \le n} \Pi_m$.

---

[5] A partition of $C \in \mathcal{E}$ is a set $\mathcal{P}$ of non-null events (that is, $A \neq 0$) that are pairwise disjoint (that is, $A \wedge B = 0$ if $A \neq B$) and such that $\bigvee_{A \in \mathcal{P}} A = C$. Partition $\mathcal{P}'$ is a refinement of partition $\mathcal{P}$ if, for each element $A \in \mathcal{P}$, a subset $\mathcal{R} \subseteq \mathcal{P}'$ is a partition of A.

3. SEU rationality

Conformity of a decision rule or of a contingent plan to subjective-expected-utility theory can be defined in terms of signed measures. A signed measure on $\mathcal{E}$ is a function $\mu\colon \mathcal{E} \to \Re$ that satisfies

$$\mu(0) = 0 \quad \text{and} \quad \forall A \ \forall B \ \left[A \wedge B = 0 \Longrightarrow \mu(A \vee B) = \mu(A) + \mu(B)\right]. \tag{5}$$

Note that in a probability space, a signed measure is defined on the field of events (that is, measurable subsets of the sample space) by considering the integral of a measurable function (from the sample space to $\Re$) as a function of the domain of integration. In particular, suppose that $(\Omega, \mathcal{E}, \Pr)$ is a probability space and that $v\colon \mathcal{A} \times \Omega$ is a state-contingent utility function. Then, for each $a \in \mathcal{A}$, a signed measure μ_a is defined by $\mu_a(B) = \int_B v(a, \omega)\, d\Pr(\omega)$. The conditional expected utility of act a in event B is $\mu(B)/\Pr(B)$. Thus a' has higher conditional expected utility than does a in event B if and only if $\mu_{a'}(B) > \mu_a(B)$.

This observation motivates a definition of SEU rationality. To state this definition, let $\mathcal{M}$ denote the set of signed measures. If $u\colon \mathcal{A} \to \mathcal{M}$, then I will write $u_a(B)$ instead of $[u(a)](B)$ in order to avoid cumbersome notation.

Let $\mathcal{C} \subseteq \mathcal{E}$, and let $\gamma\colon \mathcal{C} \to \mathcal{A}$. Then γ is *SEU rational* if there is a mapping $u\colon \mathcal{A} \to \mathcal{M}$ such that

$$\forall C \in \mathcal{C} \quad \gamma(C) = \left\{a \middle| \forall a' \ \ u_{a'}(C) \leq u_a(C)\right\}. \tag{6}$$

4. A decision rule that is not SEU rational

The following is an example of a decision rule that is not SEU rational. Let $\mathcal{E}$ be the field of subsets of a set of four elements. Specifically, let

$$1 = \{a, b, c, d\}; \ \ \mathcal{E} = \{B | B \subseteq S\}. \tag{7}$$

For an event $B \in \mathcal{E}$, define $\#B$ to be the cardinality of B. Also let

$$\mathcal{A} = 1 \cup \{f, g\}. \tag{8}$$

Define $\delta\colon \mathcal{E} \to \mathcal{A}$ by

$$\delta(B) = \begin{cases} \mathcal{A}, & \text{if } B = \emptyset; \\ B, & \text{if } B = \{x\}; \\ \{f\}, & \text{if } \#B = 2 \text{ or } B = 1; \\ \{g\}, & \text{if } \#B = 3. \end{cases} \tag{9}$$

That δ is not SEU rational is easily demonstrated by contradiction. Let $A = \{a\}$, $J = \{b, c\}$, and $K = \{b, c, d\}$,. Suppose that the mapping $u\colon \mathcal{A} \to \mathcal{M}$ were to satisfy the SEU-rationality condition (6). By (6) and (9), $u_f(J) > u_g(J)$ and $u_f(A \vee J) < u_g(A \vee J)$. By the additivity property (5) of a signed measure, therefore, $u_f(A) < u_g(A)$. However, analogous reasoning involving substitution of K for J and transposition of f and g, leads to the contradictory conclusion that $u_g(A) < u_f(A)$.

It is well known that there exist decision rules that are not SEU rational. What makes this example noteworthy is the additional fact, to be proved below, that every contingent plan induced by δ (and an arbitrary sequence of partitions satisfying (4)) does conform to SEU theory. To provide some intuition for that result, let me explain now why the demonstration that δ is not SEU rational cannot be applied to any contingent plan induced by δ. Recall that the domain of a contingent plan π is the union of a sequence of partitions ordered by refinement. It is easily seen that, if A and B are any two elements of such a union, then either $A \wedge B = A$ or $A \wedge B = B$ or $A \wedge B = \emptyset$. (This can be proved by induction on the number of partitions in the sequence, and it is also intuitively obvious from drawing a Venn diagram of such a sequence.) In the demonstration that δ is not SEU rational, information about the images of δ at both $A \vee J$ and K are used. Observe that $(A \vee J) \wedge K = J$, and that $J \neq A \vee J$ and $J \neq K$, so that $A \vee J$ and K cannot both occur in a sequence of partitions ordered by refinement. In any contingent plan, then, the decision maker's choice at one of these events or the other will be unobserved.

The kind of example presented here is pervasive. To see this, extend the cardinality operator # to events in an arbitrary finite Boolean algebra by specifying that $\#B$ is the cardinality (in the usual sense) of the finest partition of B. (Thus $\#1$ denotes the "cardinality" of the unit element of $\mathcal{E}$. If $\mathcal{E}$ is the field of all subsets of a set S, then $1 = S$, so $\#1$ is the cardinality of S.) Also, if $A \in \mathcal{E}$, then define $\mathcal{E}_A$ to be the Boolean algebra of elements $A \wedge B$ for all $B \in \mathcal{E}$. (In particular, $\mathcal{E}_1 = \mathcal{E}$.) The following lemma, which is proved in the same manner as the analysis that has just been made of the example, shows how to construct a decision rule that violates SEU rationality at many places.

Lemma 1: Let $\mathcal{E}$ be a finite Boolean algebra. Define $\mathcal{C} = \{C \in \mathcal{E} | \#C \text{ is odd and } \#C \leq (\#1)/2\}$. Let $\mathcal{A} = \mathcal{C} \cup \{f, g\}$, where $f \neq g$ and $\mathcal{C} \cap \{f, g\} = \emptyset$. Define $\delta \colon \mathcal{E} \to \mathcal{A}$ by

$$\delta(B) = \begin{cases} \mathcal{A}, & \text{if } B = \emptyset; \\ \{B\}, & \text{if } B \in \mathcal{C}; \\ \{f\}, & \text{if } \#B \text{ is even}; \\ \{g\}, & \text{if } \#B \text{ is odd and } \#B > (\#1)/2. \end{cases} \tag{10}$$

Then, for any $A \in \mathcal{E}$ such that $\#A \geq (\#1)/2 + 3$, $\mathcal{E}_A$ is not SEU rational.

5. SEU-rational contingent plans

Although the decision rules constructed in the preceding section are not SEU rational, they always induce SEU-rational contingent plans. This fact will be proved by means of the following theorem, which adapts to the present context a result of Green and Park (1993). The theorem is proved in the appendix.

Theorem 1: A contingent plan π with domain $\bigcup_{m \leq n} \Pi_m$ is SEU rational if it satisfies three conditions for all events A, B, and C in its domain.

$$\text{If } A = B \vee C \text{ and } B \wedge C = 0 \text{ and } a \in \pi(B) \text{ and } a \in \pi(C)\text{, then } a \in \pi(A). \tag{11}$$

$$\text{If } A = B \vee C \text{ and } B \wedge C = 0 \text{ and } a \in \pi(B) \text{ and } a \in \pi(C) \text{ and } b \notin \pi(B)\text{, then } b \notin \pi(A). \tag{12}$$

$$\pi(A) \neq \emptyset. \tag{13}$$

Condition (11) states that, if event A is partitioned into events B abd C, and if the decision maker would be willing to choose act a if he were certain of either B or C, then he must be willing to chose a if he is certain of event A. Condition (12) states that if the decision maker has at least one such "sure-thing" act, and if there is one of the two specific events comprised by A in which he would be unwilling to choose another act b, then he must not be willing to choose b in event A. Condition (13), which states that in any event there must be at least one act that the decision maker would be willing to choose, is motivated by the idea that the decision maker's conditional choices reflect optimization with respect to consistent preferences on the finite set $\mathcal{A}$.

Green and Park (1993) use this theorem provide some idea of the stringency of conditions (11)—(13). They show that a single-valued minimax-loss or minimax-regret decision rule always (that is, for any partition sequence) induces a contingent plan that is SEU-rational. That result does not address the specific question being studied in this paper, though, because it is consistent with the possibility that a minimax-loss or minimax-regret decision rule may always be SEU rational on its full domain if it is single-valued. However, Green and Park do address the present question by constructing an example in which a decision rule induces a contingent plan that is not SEU rational. The decision rule in the example is obtained by maximizing an instance of Chew's (1983) "weighted-utility" preference. It is noteworthy that Chew introduced that class of preferences in order to provide a relatively parsimonious generalization of expected-utility theory that would accomodate experimental evidence regarding Allais' paradox. Thus the example would support the view that behavior that is inconsistent with SEU rationality should presumably be observable outside the laboratory, if it can be elicited from experimental subjects inside the laboratory. Now I turn to the analysis of the examples constructed in the proceding section, and I show that such a view does not necessarily have to be taken.

6. SEU rationality of contingent plans induced by SEU-irrational decision rules

The decision rules defined at the beginning of section 4 and in lemma 1 can be shown always to induce SEU-rational contingent plans, by showing that it satisfies conditions (11)—(13) of theorem 1.[6] Specifically, consider the decision rule δ defined by condition (10) in lemma 1. Clearly δ satisfies condition (13) of being nonempty valued. Moreover, since δ is single valued (except at 0), condition (11) implies condition (12). To see this, suppose that the antecedent of (12) is true. Then, by (11), $a \in \delta(A)$. Since $b \notin \pi(B)$, $B \neq 0$ and therefore $A \neq 0$. Since δ is single valued at A, $b \notin \delta(A)$. That is the conclusion of (12), so the implication (12) is true.

Thus only condition (11) has to be checked. Suppose, therefore, that $A = B \vee C$ and that $B \wedge C = 0$. To consider the case that $B = 0$ or $C = 0$, assume without loss of generality that $B = 0$. Then $A = C$, so (11) is tautological. If neither $B = 0$ nor $C = 0$, then there are three possibilities. Either both $\#B$ and $\#C$ are even, or else both $\#B$ and $\#C$ are odd and at least one of them is no larger than $(\#1)/2$ (since $B \wedge C = 0$), or else one of $\#B$ and $\#C$ is even and the other is odd. In the first case, $\#A$ is also even, so $\delta(B) = \delta(C) = \delta(A) = \{f\}$ by definition (10), so (11) is true. In both of the remaining cases, (10) implies that there are x and y such that $\delta(B) = \{x\}$ and $\delta(C) = \{y\}$ and $x \neq y$, so the antecedent of implication (11) must false because there can be no a such that $a \in \delta(B)$ and $a \in \delta(C)$. Therefore (11) is true. This proves the following result.

Theorem 2: If π is induced by the decision rule δ defined by (10), together with any sequence of partitions satisfying (4), then π is SEU rational.

---

[6] Lemma 1 does not show that the decision rule in the example given at the beginning of section 4 is not SEU rational, because in that example $\#1 < (\#1)/2 + 3$. The bound $(\#1)/2 + 3$ can actually be improved to $(\#1)/2 + 2$, though, except in the case that $\#1$ is divisible by 2 but not by 4. That improvement subsumes the example.

Appendix: Proof of Theorem 1

Conditions (11) and (12) imply that

$$\text{If } \Theta \text{ is a partition of } \mathcal{A}, \text{ then } \bigcap_{B\in\Theta} \pi(B) = \pi(A) \text{ or } \bigcap_{B\in\Theta} \pi(B) = \emptyset. \tag{14}$$

Signed measures u_a for $a \in \mathcal{A}$ will be defined from a function $V{:}\, \mathcal{A} \times \bigcup_{m\leq n} \Pi_m \to \Re$. The function V is defined recursively on the partitions Π_m, as follows. The basis step is that, if $B \in \Pi_0$, then $B = 1$ and

$$V(a,1) = \begin{cases} 1, & \text{if } a \in \pi(1); \\ 0, & \text{if } a \notin \pi(1); \end{cases} \tag{15}$$

For the recursion step, let $\phi{:}\, \bigcup_{m\leq n} \Pi_m \to \mathcal{A}$ be a selection from π. (That is, $\forall A \in \bigcup_{m\leq n} \Pi_m \;\; \phi(A) \in \pi(A)$.) Suppose that $m < n$ and that V has been defined on $\mathcal{A} \times \Pi_m$. The next partition Π_{m+1} can be expressed uniquely as a union $\bigcup_{C\in\Pi_m} \Theta_C$, where each Θ_C is a partition of C. For each $a \in \mathcal{A}$ and for each $C \in \Pi_m$, define Θ_{aC} to be the set of $B \in \Theta_C$ such that either $a \in \pi(B)$ or else $\forall D \in \Theta_C \;\; a \notin \pi(D)$. Now, for each $a \in \mathcal{A}$, extend the definition of V to $\{a\} \times \Theta_{aC}$ by specifying that

$$V(a,B) = \begin{cases} V(\phi(C),C), & \text{if } a \in \pi(B) \text{ and } \bigcap_{B\in\Theta_C} \pi(B) \neq \emptyset; \\ V(\phi(C),C) + 1, & \text{if } a \in \pi(B) \text{ and } \bigcap_{B\in\Theta_C} \pi(B) = \emptyset; \\ V(a,A), & \text{if } \forall D \in \Theta_C \;\; a \notin \pi(D); \end{cases} \tag{16}$$

Complete the extension of V to $\mathcal{A} \times \Theta_C$ by specifying that

$$\text{If } B \notin \Theta_{aC}, \text{ then } V(a,B) = \rho < V(a,C), \tag{17}$$

where ρ solves the equation

$$\sum_{D\in\Theta_{aC}} \left[V(a,D) \cdot \#D\right] + \rho\left[\#C - \sum_{D\in\Theta_{aC}} \#D\right] \; = \; V(a,C) \cdot \#C. \tag{18}$$

It must be proved that the equation (18) always possesses a solution, and that this solution satisfies the inequality in (17). Both of these facts are guaranteed by (14), as is shown by Green and Park (1993).

Conditions (17) and (18) are essentially a martingale condition, with respect to counting measure on the atoms of $\mathcal{E}$ (that is, the elements $D \in \mathcal{E}$ such that $\#D = 1$), for the conditional expected utility $V(a,B)$ of taking act a. By the martingale convergence theorem, for each a there is a signed measure $u_a \in \mathcal{M}$ that satisfies

$$\forall B \in \bigcup_{m\leq n} \Pi_m \quad u_a(B) = V(a,B). \tag{19}$$

By induction on m, using conditions (16) and (17),

$$\forall B \in \bigcup_{m\leq n} \Pi_m \quad \pi(B) = \{a \mid \forall a' \in \mathcal{A} \;\; V(a',B) \leq V(a,B). \tag{20}$$

Conditions (19) and (20) together imply condition (6) which defines SEU rationality. ∎

References

Chew, S.-H., "…quasilinear-mean utility…," *Econometrica,* 1983.

Green, E. J., and I.-U. Park, "Bayes contingent plans," w.p. 9307002, ewp-game@econwpa.wustl.edu.

Hogarth, R. M., and M. W. Reder, "Rational Choice," University of Chicago Press, 1987. (Previously published as a supplement to the *Journal of Business,* 1986.)

Rust, J., "Optimal replacement of GMC bus engines: an empirical study of Harold Zurcher," *Econometrica,* 1987.