The Effect of Exchange Rates on Statistical Decisions

Mark J. Schervish, Teddy Seidenfeld,
and Joseph B. Kadane*†

Statistical decision theory, whether based on Bayesian principles or other concepts such as minimax or admissibility, relies on minimizing expected loss or maximizing expected utility. Loss and utility functions are generally treated as unit-less numerical measures of value for consequences. Here, we address the issue of the units in which loss and utility are settled and the implications that those units have on the rankings of potential decisions. When multiple currencies are available for paying the loss, one must take explicit account of which currency is used as well as the exchange rates between the various available currencies.

1. Introduction. Normative decision theory, and in particular statistical decision theory, is generally based on minimizing a quantitative loss function or maximizing a quantitative utility function, which is itself a unit-less representation of values for outcomes. Such representations result from the familiar axiomatic derivations of subjective expected utility theory (see, e.g., Savage 1954; Anscombe and Aumann 1963). There is also minimax theory (Wald 1950), wherein the decision maker chooses an option so as minimize, over all feasible options, the maximum, over states, of the loss. Minimax theory and others assume the existence of a loss function without addressing the issue of in what, if any, units the loss is measured.

What are the effects of not being careful about the units in which losses and gains are realized when outcomes are in terms of one set of prizes rather

Received October 2012; revised May 2013.

*To contact the authors, please write to: Teddy Seidenfeld, 135J Baker Hall, Pittsburgh, PA 15213; e-mail: teddy@stat.cmu.edu.

†The authors thank two anonymous referees for their helpful suggestions in preparing this article. The first author was supported, in part, by a grant from the Institute of New Economic Thinking.

Philosophy of Science, 80 (October 2013) pp. 504–532. 0031-8248/2013/8004-0004$10.00
Copyright 2013 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/673267 Published online by Cambridge University Press

than another? Example 1 is a toy example that will be extended to illustrate a problem that arises whenever one is not careful about how a loss function is actually paid. The example and its extensions will raise serious issues that need to be addressed about whether the loss function is measured in pure numbers or in units with intrinsic values.

Although we focus primarily on statistical decision theory in this article, the issues raised by example 1 and its extensions have wide-ranging consequences in philosophy whenever uncertainty is connected with decision making. For instance, the issues affect how the criterion of *coherence* (avoiding sure loss when betting) regulates rational degrees of belief (Shimony 1955). The same issues affect how bargaining theory (e.g., the Nash bargaining solution) may regulate ethical standards of fairness (Braithwaite 1955; Harsanyi 1977).

EXAMPLE 1. Let $\{E, E^C\}$ be a partition into two complementary events, E and not-E (denoted E^C), sufficient for formulating the following decision problems. Let $\mathcal{D} = \{d_1, d_2\}$ be the set of feasible options with a loss function given by the following table:

	E	E^C
d_1	.5	2.0
d_2	1.0	1.0

The minimax decision in this problem would appear to be to choose d_2 since the maximum loss is 1, whereas the maximum loss for choosing d_1 is 2. Also, a naive Bayesian who uses a prior distribution with $\Pr(E) = \Pr(E^C) = 0.5$ will make the same decision since the expected loss for choosing d_1 is 1.25, while the expected loss for choosing d_2 is 1.

In this article, we take seriously the question, "In what units will the agent in example 1 (and every other decision problem with a loss function) be charged the stated loss?" There are two natural answers to this question, both of which will be illustrated in versions of example 1, and neither of which is handled in a satisfactory fashion by the popular methods used in statistical decision theory.

Two natural assumptions about the units in which an agent pays a loss function are:

i) The agent must pay an amount of some commodity or currency of value equal to the numerical value of the loss function.
ii) For each combination of decision and state of nature, the agent must pay a specified amount of some commodity or currency. The agent has a cardinal utility function over amounts of said commodity or currency, and the numerical values of the loss function encode the agent's loss of utility from losing the specified amount of com-

https://doi.org/10.1086/673267 Published online by Cambridge University Press

modity/currency in each of the combinations of decision/state of nature.

There is a third assumption that we will not pursue in this article, namely, that no loss is paid by the agent, and the loss function is just a mathematical construct with no operational meaning.

Assumption ii is a generalization of i, in the sense that the two are essentially the same if the agent's utility is linear in the units of the commodity/currency. The problems that arise for both interpretations are easily illustrated by extensions of example 1. First, consider assumption i.

EXAMPLE 2. In example 1, suppose that the loss is paid in a currency C_1, while there is an alternative currency C_2 whose exchange rate with C_1 is given as follows:

- On event E, one unit of C_1 is worth two units of C_2.
- On event E^C, one unit of C_1 is worth 0.5 units of C_2.

Suppose that the minimax decision maker prefers to think in units of C_2. The loss function from example 1, converted to units of C_2, is

	E	E^C
d_1	1.0	1.0
d_2	2.0	.5

The minimax decision is now d_1, even though the two loss functions charge equivalent values under all circumstances. Application of the minimax rule regardless of the units in which losses occur results in inconsistent decision making. If the naive Bayesian in example 1 calculates expected loss independent of the currency, he or she will make the same choices as the minimax decision maker. With currency C_2, the expected losses for decisions d_1 and d_2 are 1.0 and 1.25, respectively.

The seemingly inconsistent decisions in examples 1 and 2 arise from a failure to account for the varying values of one unit of loss from state to state. The minimax theory explicitly treats one unit of loss as being equally important in every state. But the state-dependent exchange rate makes it clear that one unit of loss cannot be equally valuable in both states for both currencies. The minimax decision maker either needs to guarantee that one unit of loss means the same thing in every state or needs explicitly to take into account the varying value of one unit of loss.

Axiomatic derivations of decision theory (e.g., Savage 1954; Anscombe and Aumann 1963) make assumption ii explicit by encoding the decision maker's values through a unit-less, real-valued utility function. These derivations derive the utility by giving privileged status to a particular set of

https://doi.org/10.1086/673267 Published online by Cambridge University Press

constant acts. These constant acts purport to serve as utility theory's "rigid rods" (or *numeraires*, in the language of mathematical finance) for measuring both preference and uncertainty across states of uncertainty. The problem with applying either the theory of Anscombe and Aumann (1963) or the theory of Savage (1954) in examples 1 and 2 is the choice of numeraire. Either currency C_1 or C_2 can play the role of numeraire, but the mere presence of the other currency violates an assumption of each of the subjective expected utility theories. This violation prevents the identification of a unique representation of an agent's preference in terms of a single probability and a single utility. The same problem arises when applying bargaining theory in the service of ethics, as in the works of Braithwaite (1955) and Harsanyi (1977). As we explain in section 2, a more general expected utility theory is needed to allow for relative values that change from event to event. Unfortunately, the more general theory makes it clear why it is not possible to compare numerical utility values from one event to the next. For example, regardless of which of the two loss functions (in examples 1 and 2, respectively) corresponds to an agent's cardinal utility, d_2 costs the agent twice as much as d_1 costs when E occurs. Similarly, d_1 costs twice as much as d_2 costs when E^C occurs in both examples. What we cannot say is whether d_1 costs more or less when E occurs than when E^C occurs. Such comparisons are not available because they are confounded with the probabilities of the events E and E^C. Another extension of example 1 will illustrate this.

> EXAMPLE 3. Suppose that the agent in examples 1 and 2 faces a collection of decision problems, including the one described in those examples. The payments that the agent will be charged will be either one unit of currency C_1 or one unit of currency C_2. Using assumption i to explain what the loss function measures, the decision problem in examples 1 and 2 corresponds to the agent paying one unit of currency C_2 if decision d_1 is made (regardless of whether E or E^C occurs) and one unit of currency C_1 if decision d_2 is made (regardless of whether E or E^C occurs). The different versions of the loss function in those two examples correspond to different choices of the numeraire. In example 1, the number 1 stands for one unit of C_1 in both events, while in example 2, the number 1 stands for one unit of C_2 in both events. Suppose that the naive Bayesian who used probability $\Pr(E) = \Pr(E^C) = 0.5$ in example 1 were to change the probability to $\Pr(E) = 0.2$ and $\Pr(E^C) = 0.8$ in example 2. The expected losses in example 2 would now be 1 for d_1 and 0.8 for d_2. Not only would the agent choose d_2 in both examples, but the ratio of the two expected losses is 0.8 in both examples.

It should come as no surprise that changing the probability used to compute expected loss might change the chosen decision. What happens in example 3 is actually characteristic of every statistical decision problem. When one changes the numeraire (i.e., what counts as one unit of loss), then there

https://doi.org/10.1086/673267 Published online by Cambridge University Press

is a corresponding change to the probability over the states of nature that will preserve all statistical decisions. We extend example 1 again to illustrate this.

EXAMPLE 4. There are a total of 16 decision problems involving two decisions, the two events (E and E^C), and the two prizes (one unit of either currency) that can be distinguished by what the agent must pay. For this example, we make assumption ii, to explain what the loss function measures, in order to show that the phenomenon of example 3 does not depend on assumption i. Denote the loss from paying one unit of currency i ($i = 1, 2$) by $L(i, E)$ or $L(i, E^C)$, depending on which event occurs. The loss function for each of the 16 decision problems corresponds to a table similar to those in examples 1 and 2. For decision d_j ($j = 1, 2$), let $i_j(F)$ denote the currency that the agent must pay (one unit) if he or she makes decision d_j and event F (either E or E^C) occurs. Then the 16 loss functions all have the form:

	E	E^C
d_1	$L(i_1(E), E)$	$L(i_1(E^C), E^C)$
d_2	$L(i_2(E), E)$	$L(i_2(E^C), E^C)$

Suppose that the agent combines the above loss function with a probability $p = \Pr(E)$ and $1 - p = \Pr(E^C)$. Then, the agent chooses d_1 if

$$\begin{aligned} &pL(i_1(E), E) + (1 - p)L(i_1(E^C), E^C) \\ &< pL(i_2(E), E) + (1 - p)L(i_2(E^C), E^C) \end{aligned} \tag{1}$$

He or she chooses d_2 if the opposite inequality holds, and either decision is optimal if the inequality becomes equality.

What we show next is that, no matter how the agent specifies $L(\cdot, \cdot)$ and p (so long as p is neither 0 nor 1), there are infinitely many alternative choices that lead to identical decisions in all 16 decision problems simultaneously. That is, (1) holds for one choice of L and p if and only if it holds for all of the other choices. Each choice of L and p corresponds to a different choice of numeraire. Let c_E and c_{E^C} be two strictly positive finite numbers. Replace L by $L'(i, F) = L(i, F)/c_F$ for $i \in \{1, 2\}$ and $F \in \{E, E^C\}$. Let $m = pc_E + (1 - p)c_{E^C}$. Replace p by $q = pc_E/m$ so that $1 - q = (1 - p)c_{E^C}/m$. It is simple to see that (1) holds if and only if

$$\begin{aligned} &qL'(i_1(E), E) + (1 - q)L'(i_1(E^C), E^C) \\ &< qL'(i_2(E), E) + (1 - q)L'(i_2(E^C), E^C). \end{aligned} \tag{2}$$

Indeed, the two sides of (2) are merely $1/m$ times the corresponding sides of (1). Example 3 corresponds to using $c_E = 0.5$, $c_{E^C} = 2$, and $p = 0.5$ to transform the loss table in example 1.

https://doi.org/10.1086/673267 Published online by Cambridge University Press

The important point to learn from example 4 is that the relative values of losses on disjoint events are hopelessly confounded with the probabilities of those events. It is meaningless to compare the numerical values of the losses that arise on different events without reference to the probabilities. We can rescale the loss function to make the ratio of losses (even for the same payment) on different events to be any positive finite value we want, without affecting what decisions an agent will make, so long as we modify the probabilities accordingly. Similarly, we can modify the probabilities of any event E (except one whose probability is either 0 or 1) so that $\Pr(E)/\Pr(E^C)$ is anything we want. None of this invalidates subjective expected utility theory, as we show in section 2. However, it is difficult to justify treating probability values (other than 0 and 1) as if they had meanings independently of the utility (or loss) function. Similarly, the utility function values do not have meaning independently of the probabilities of events. Rubin (1987) raised this issue, but exploration of its implications has been limited in the years since. This interaction between the state-dependent value of prizes and a rational agent's degrees of belief also affects the traditional Dutch Book betting arguments that equate a decision maker's coherent fair-betting rates with the decision maker's confirmation function (Shimony 1955).

The examples in this section are simplified decision problems in which some subtle issues have been avoided:

1. How shall we formalize a global utility function that applies simultaneously across different numeraires?
2. What may serve as a numeraire for formulating gambles, and how is the exchange rate between numeraires related to the decision maker's uncertainty over states?
3. Can preferences be elicited by special loss functions, for example, using scoring rules, in order to identify the decision maker's quantitative degrees of belief—where degrees of belief are epistemic states that are not indexed by the units in which prizes are given?

In this article, we provide a general theory to answer these questions and others that we identify later on. However, our positive approach does not salvage the inconsistent decisions generated by the minimax theory or by the naive Bayesian's analysis, where the decision maker's uncertainty is represented by a probability that ignores the units associated with prizes.

In sections 2 and 3, we address the first two questions in considerable generality. We show how a general theory of maximizing expected utility takes explicit account of the conversion from currencies and commodities of value into pure utility values. We focus on the state-dependent relative values between various currencies and commodities to which attention must be paid in order to avoid inconsistencies like those displayed in examples 1

https://doi.org/10.1086/673267 Published online by Cambridge University Press

and 2. Also, in many applications, utility functions are not linear in every currency. We address this in section 4.

Under the theory developed here, in general there is no unique subjective probability that the Bayesian can use in all decision problems without regard to how losses are paid. However, as considered in section 5, we give conditions under which the same decision will be made, regardless of which of several currencies is used for paying the loss. In this connection, in section 6 we investigate elicitation and question 3 (above). We consider special decision problems in which a decision maker is asked to provide a subjective expected value for some random variable. For instance, the decision maker's degree of belief about an event E is his or her expected value of the indicator function for E. This is the area in which the state-dependent theory has its most striking consequences. The idea common in the work of De Finetti (1974) on Brier score and the more general analysis of proper and consistent scoring rules of Savage (1971) and Gneiting (2011a, 2011b) is that one may elicit a decision maker's degrees of belief about specific random quantities (i.e., the indicator function for E) by providing the decision maker with incentive-compatible losses for the elicitation. However, as we show, this approach needs to be tempered by the realization that the elicited probability is just one of many that form part of a state-dependent expected utility representation of the same preference ordering. We show how the choice of unit for applying the scoring rule affects the elicitation. That is, we show where state-dependent utilities create strategic aspects for elicitation, even with proper/consistent scoring rules.

2. State-Dependent Utility. Let Ω be a set of *states of nature*, that is, any partition of the sure event. In a typical mathematical presentation, Ω would have a σ-field $\mathcal{A}$ of subsets. Measurable real-valued functions defined on Ω are called *random variables*. Elements of $\mathcal{A}$ are called *events*, and we allow ourselves the convention of denoting the indicator function of an event $\mathcal{A}$ by the name of the event itself. If P is a probability on $(\Omega, \mathcal{A})$ and X is a random variable, we will allow ourselves the convention of letting $P(X)$ stand for the expected value of X under P, $\int_\Omega X(\omega)\, dP(\omega)$.

Let $\mathcal{R}$ be a set of fortunes for a decision maker with σ-field $\mathcal{B}$ of subsets. A Von Neumann–Morgenstern lottery (NM lottery) L is a stipulated probability distribution (auxiliary randomization) over the set $\mathcal{R}$. Let $\mathcal{H}$ be a set of functions from Ω to the NM lotteries. (See Von Neumann and Morgenstern [1947] for a discussion of how NM lotteries figure in the axiomatic derivation of decision theory.) An element H of $\mathcal{H}$ is called a *horse lottery*, following Anscombe and Aumann (1963). There is one special element of $\mathcal{R}$ that we will call *status quo*. It stands for the current fortune of a decision maker at the point when he or she is being asked to make the next decision.

https://doi.org/10.1086/673267 Published online by Cambridge University Press

We assume that, in every state ω there is some fortune better than status quo and some fortune worse than status quo.

Anscombe and Aumann (1963) prove that an agent's preferences among simple horse lotteries satisfy some seemingly innocuous conditions, if and only if they can be represented by a unique probability/utility pair. That is, the conditions hold if and only if there is a unique probability P and a unique utility, a bounded function $U\colon \mathcal{R} \to \mathbb{R}$, with the following property. The agent prefers H_2 to H_1 if and only if

$$P[U(H_1)] < P[U(H_2)]. \tag{3}$$

When L is an NM lottery, the meaning of $U(L)$ is $\int_R U(v)\, dL(v)$. Savage (1954) gives an alternative derivation of an expected-utility representation of preference. See Fisburn (1970) for an overview of several derivations of expected utility theory.

One of the seemingly innocuous conditions of Anscombe and Aumann (1963) implies that, almost surely, the relative values of fortunes remain the same as the state of nature changes (*state independence*). This condition is violated when the fortunes involve different currencies whose values can vary from state to state with positive probability. Without that state-independence condition, the uniqueness of the probability/utility representation is lost, and the utility function must be a more general object of the form to be defined in definition 1 below.

> EXAMPLE 5. Let $\Omega = \{1, 2\}$. Suppose that \$1 can be exchanged cost-free for €0.7 in state 1 and that \$1 can be exchanged cost-free for €0.6 in state 2. Suppose also that utility is continuous and increasing in both currencies but is not necessarily linear in either currency. Finally, suppose that a state-independent utility representation for preference, as in (3), gives each state probability 0.5. Let $U(€0.6) = a$, $U(€0.7) = b$, and $U(\$1) = c$. Then $c = 0.5a + 0.5b$, and $a < c < b$. Let x be the number of euros that has value c; that is, $U(€x) = c$. Then $0.6 < x < 0.7$. The state-independent utility representation would assign expected utility c to the horse lottery H that gives \$1 in state 1 and €x in state 2. But this is unsatisfactory since H is the same as \$1 in state 1 and H is strictly more valuable than \$1 in state 2, which has positive probability.

State-independent utility representations such as (3) are not capable of representing preferences when the relative values of fortunes vary from state to state. Hence, we introduce the usual generalization to handle such cases. (See Rubin [1987] for one derivation.)

https://doi.org/10.1086/673267 Published online by Cambridge University Press

DEFINITION 1 (State-Dependent Utility). Let $\mathcal{P}$ be a collection of mutually absolutely continuous probabilities on Ω. Suppose that, for each $P \in \mathcal{P}$, there is a utility function U_P: $\Omega \times R \to \mathbb{R}$ with the following properties.

- For all P,

$$P[\sup_v |U_P(\cdot, v)|] < \infty. \tag{4}$$

- For every $P_1, P_2 \in \mathcal{P}$,

$$U_{P_2}(\omega, v) = c_{1,2} U_{P_1}(\omega, v) \frac{dP_1}{dP_2}(\omega) + t_{1,2}(\omega), \tag{5}$$

almost surely, where $c_{1,2} > 0$ is a scalar that can depend on P_1 and P_2 but nothing else, $t_{1,2}$ is some P_2-integrable function of ω, and dP_1/dP_2 is the Radon-Nikodym derivative of P_1 with respect to P_2.

The collection $\{(P, U_P): P \in \mathcal{P}\}$ is called a *state-dependent expected utility representation* of preference over $\mathcal{H}$. We say that a horse lottery H has state-independent values under U_P if $U_P(\omega, H(\omega))$ is constant as a function of ω.

EXAMPLE 6. Every state-independent expected utility representation of preference extends in a simple fashion to a state-dependent utility. Let P be a probability and let U be a bounded function such that the agent prefers H_2 to H_1 if and only if (3). Let $\mathcal{P}$ consist of all probabilities that are mutually absolutely continuous with respect to P. For each $Q \in \mathcal{P}$, define

$$U_Q(\omega, v) = U(v) \frac{dP}{dQ}(\omega).$$

It is straightforward to see that $\{(Q, U_Q): Q \in \mathcal{P}\}$ satisfies definition 1 with $c_{1,2} = 1$ and $t_{1,2}(\cdot) \equiv 0$.

Throughout this article, we assume that $\mathcal{P}$ is as large as possible in the following sense. If $P_1 \in \mathcal{P}$ and P_2 is mutually absolutely continuous with P_1, then $P_2 \in \mathcal{P}$. This causes no loss of generality because U_{P_2} is easily constructed from (5). We will also suppress the "almost surely" qualification in equations and formulas that involve Radon-Nikodym derivatives since all probabilities in $\mathcal{P}$ have the same zero-probability sets.

Condition (4) in definition 1 is the state-dependent analog of the requirement that a utility is a bounded function. Even if some U_P are bounded, the conversion equation (5) allows other U_P to be unbounded. Because each utility can be multiplied by a positive constant without changing the repre-

https://doi.org/10.1086/673267 Published online by Cambridge University Press

sentation of preference, (4) means that all expected utilities can be bounded by a common bound. Technically, the condition that utilities are bounded arises as a consequence of such derivations as Savage (1954). If one skips the derivation of expected utility and simply adopts a probability/utility pair (P, U) as in example 6, one need not assume that U is bounded, so long as one can guarantee that the expected utilities of all horse lotteries are finite. This would require restrictions on the set $\mathcal{H}$ of horse lotteries.

EXAMPLE 7. In example 6, assume that the utility U is unbounded above. For each n, let v_n be a fortune such that $U(v_n) > 2^n$. If L_0 is the NM lottery that assigns fortune v_n with probability 2^{-n}, then $U(L_0) = \infty$. If $P[H = L_0] > 0$, then $P[U(H)]$ will be either infinite or undefined. Clearly, we cannot allow elements of $\mathcal{H}$ to assume NM lotteries like L_0. For infinite spaces, we need further restrictions on $\mathcal{H}$. Assume that there are disjoint subsets $\{A_n\}_{n=1}^{\infty}$ of Ω such that $P(A_n) = a_n > 0$ for all n. For each n, let w_n be a fortune such that $U(w_n) > 1/a_n$. Let $H_0 = \sum_{n=1}^{\infty} A_n w_n$. That is, for each n and each $\omega \in A_n$, $H_0(\omega) = w_n$. Then $P[U(H_0)] > \sum_{n=1}^{m} U(w_n)a_n > m$ for every natural number m. Hence, $P[U(H_0)] = \infty$. In order for $P[U(H)] < \infty$ for all $H \in \mathcal{H}$, we must prevent H_0 and all similar horse lotteries from being in $\mathcal{H}$. One way to do that would be to restrict $\mathcal{H}$ to contain only simple horse lotteries, namely, those that assume only finitely many NM lotteries, each of which has finite utility, as in Anscombe and Aumann (1963).

Rather than impose the types of restrictions discussed in example 7, we assume (4). Seidenfeld, Schervish, and Kadane (2009) discuss other problems that arise when utilities are unbounded.

One important consequence of (5) is as follows. Let H_1, H_2 be elements of $\mathcal{H}$. Then, for every $P_1, P_2 \in \mathcal{P}$,

$$\begin{aligned} &\int_{\Omega} U_{P_1}(\omega, H_1(\omega))\, dP_1(\omega) < \int_{\Omega} U_{P_1}(\omega, H_2(\omega))\, dP_1(\omega) \\ &\qquad\qquad \text{if and only if} \\ &\int_{\Omega} U_{P_2}(\omega, H_1(\omega))\, dP_2(\omega) < \int_{\Omega} U_{P_2}(\omega, H_2(\omega))\, dP_2(\omega). \end{aligned} \tag{6}$$

That is, every probability/utility pair (P, U_P) ranks all horse lotteries the same as every other such pair.

It is easy to see that one can add an arbitrary integrable function of ω to a utility or multiply a utility by a positive constant without changing how the utility ranks horse lotteries. We will make a standardization of all utility

https://doi.org/10.1086/673267 Published online by Cambridge University Press

functions so that $U_P(\omega, \text{ status quo}) = 0$, for all ω and all P. Hence, status quo has the state-independent value 0 under all utilities. In (5), this makes $t_{1,2}$ identically 0 for all P_1 and P_2.

The scalar factor $c_{1,2}$ in (5) is an inconvenience that we can do without if we scale all utilities in a standard way. There are uncountably many ways that we could scale. The most convenient way is to pick a single P_0 and force $U_P = U_{P_0} \times (dP_0/dP)$ for all other $P \in \mathcal{P}$. No matter which P_0 we choose for this purpose, we get $c_{1,2} = 1$ in (5) for all P_1 and P_2. Also, we still have $U_P(\omega, \text{ status quo}) = 0$ for all P.

With the standardizations above, we see that, for all $v \in \mathcal{R}$, (5) gives

$$U_{P_2}(\omega, v) = U_{P_1}(\omega, v) \frac{dP_1}{dP_2}(\omega), \tag{7}$$

for all P_1 and P_2. In particular for each ω and v, the sign of $U_P(\omega, v)$ is the same for all P. Also, for every horse lottery H and all $P_1, P_2 \in \mathcal{P}$,

$$\int_\Omega U_{P_1}(\omega, H(\omega))\, dP_1(\omega) = \int_\Omega U_{P_2}(\omega, H(\omega))\, dP_2(\omega). \tag{8}$$

This is a more convenient (and seemingly stronger) form of (6).

We make heavy use of a special kind of horse lottery in the rest of this article.

DEFINITION 2 (Numeraire). A numeraire is any horse lottery H such that $U_P(\omega, H(\omega))$ has the same sign (not 0) for all P and all ω. If that sign is positive, the numeraire is called *positive*, and if the sign is negative, the numeraire is called *negative*. The *marginal value* of a numeraire H is the number

$$c_H = \int_\Omega U_P(\omega, H(\omega))\, dP(\omega), \tag{9}$$

which is the same for all P according to (8).

The name "numeraire" is commonly used in finance to refer to a currency that counts as a unit for various calculations. In section 3.2, numeraires will provide a convenient stand-in for currency values when utilities are nonlinear.

LEMMA 1. Let H be a numeraire. Then there is a unique probability/utility pair (Q, U_Q) such that H has state-independent value c_H under U_Q.

Proof. Let (P, U_P) be a probability/utility pair. Let Q be the probability with $dQ/dP = U_P(\cdot, H)/c_H$. It follows from (7) that $U_Q(\omega, H(\omega)) = c_H$ for all ω, and H has state-independent values under U_Q. If $(Q', U_{Q'})$ is

https://doi.org/10.1086/673267 Published online by Cambridge University Press

another probability/utility pair for which H has state-independent values, then dQ/dQ' is constant by (7), and that constant must be 1. So $Q = Q'$ and $U_Q = U_{Q'}$. QED

DEFINITION 3. For each numeraire H, we refer to the pair (Q, U_Q) such that H has state-independent values under U_Q as the *probability and utility corresponding to H*.

Although each numeraire has state-independent values under one and only one utility, each utility may have several numeraires that all have state-independent values. For example, if \$1 has state-independent values under a utility, and if that utility is linear in dollar values, then \$2 will have state-independent values as well. With a general utility, if H has state-independent values and $0 < \alpha < 1$, then the numeraire that gives, in each state ω, $H(\omega)$ with probability α and status quo with probability $1 - \alpha$ also has state-independent values.

Lemma 1 gives some insight into how to fix the decision making in examples 1 and 2.

EXAMPLE 8. Reconsider example 1. One unit of currency C_1 is a numeraire as is one unit of C_2. They do not have the same corresponding probability/utility pairs, however. A Bayesian who uses $\Pr(\omega = 1) = \Pr(\omega = 2) = 0.5$ with one of the two currencies cannot use that same probability with the other currency. The theory does not allow it. Once we introduce general currencies and exchange rates, we can be more specific about the probabilities that correspond to the two currencies in this example.

The minimax decision maker in examples 1 and 2 behaves as if both numeraires have state-independent values, but that is impossible. If one of the numeraires has state-independent values, the other does not. A minimax decision maker needs some way to figure out which numeraire, if either, has state-independent values. Next, we turn to the general concept of currency and how it is related to utility in a state-dependent utility representation of preference.

3. Currency. We give a general definition of currencies, so that we can make precise the dependence of statistical decisions on currency.

DEFINITION 4. A *currency* is a set C of horse lotteries in one-to-one correspondence with a subset R_C of the reals $(A_C{:}C \leftrightarrow R_C)$ that satisfies the following conditions.

- R_C contains 0.
- $A_C(H_1) < A_C(H_2)$ if and only if, for every ω and every utility U and every $H_1, H_2 \in C$, $U(\omega, H_1(\omega)) < U(\omega, H_2(\omega))$.
- $A_C^{-1}(0)$ is status quo.

https://doi.org/10.1086/673267 Published online by Cambridge University Press

Currencies are defined as changes relative to the status quo and in such a way that more is always better. The reason for allowing R_C to be a subset of the reals (rather than requiring it to be the whole set of reals) is primarily as follows. In order for utility to be bounded when utility is also linear in currency, we need the set of currency values to be bounded. Definition 5 makes precise what we mean to say, that utility is linear in a currency.

DEFINITION 5. We say that utility is linear in currency C if, for each P, there exists $W_{P,C}:\Omega \to \mathbb{R}^+$ such that

$$U_P(\omega, A_C^{-1}(x)) = W_{P,C}(\omega)x, \tag{10}$$

for all ω and all $x \in R_C$. Let $\mathcal{C}$ stand for the class of all currencies C such that utility is linear in C.

Lemma 2, below, shows that (10) holds for a single $P = P_0$ if and only if it holds for all P with

$$W_{P,C}(\omega) = W_{P_0,C}(\omega)\frac{dP_0}{dP}(\omega). \tag{11}$$

3.1. General Results. The first result merely says that currency values are numeraires, and its proof is straightforward.

PROPOSITION 1. If C is a currency, then every element H of C except status quo is a numeraire with sign equal to the sign of $A_C(H)$.

The next result is useful when we try to define exchange rates. It says that the state-dependent relative values of two numeraires do not depend on the particular probability/utility pair used to represent preference.

LEMMA 2. Let H_1 and H_2 be two numeraires. Then, $U_P(\omega, H_2(\omega))/U_P(\omega, H_1(\omega))$ is the same for all $P \in \mathcal{P}$.

Proof. Let P_1 and P_2 be arbitrary probabilities in $\mathcal{P}$. It follows from (7) that

$$\frac{U_{P_2}(\omega, H(\omega))}{U_{P_1}(\omega, H(\omega))} = \frac{dP_1}{dP_2}(\omega), \tag{12}$$

for each numeraire H and for all P_1, P_2, and ω. Hence, the ratio on the left side of (12) does not depend on H. That is, for all P_1, P_2, H_1, H_2, and ω,

$$\frac{U_{P_2}(\omega, H_1(\omega))}{U_{P_1}(\omega, H_1(\omega))} = \frac{U_{P_2}(\omega, H_2(\omega))}{U_{P_1}(\omega, H_2(\omega))}.$$

https://doi.org/10.1086/673267 Published online by Cambridge University Press

Rearranging terms gives

$$\frac{U_{P_1}(\omega, H_2(\omega))}{U_{P_1}(\omega, H_1(\omega))} = \frac{U_{P_2}(\omega, H_2(\omega))}{U_{P_2}(\omega, H_1(\omega))}. \tag{13}$$

QED

3.2. Utility Linear in Currency. The next result exhibits a useful relationship between values of a currency that has linear utility values.

LEMMA 3. Let $C \in \mathcal{C}$, and for each $x \neq 0$, let $H_{C,x} = A_C^{-1}(x)$, that is, x units of currency C. Then, the probability/utility pair (P_x, U_{P_x}) corresponding to $H_{C,x}$ is the same for all $x \neq 0$, and the state-independent value of x units of currency C is $xc_{H_{C,1}}$.

Proof. Let $x \neq 0$. Because $H_{C,x}$ has state-independent values under U_{P_x}, $W_{P_x,C}$ is constant. Let $P_0 \in \mathcal{P}$. From (10) and (11), we see that P_x must satisfy

$$\frac{dP_x}{dP_0} = \frac{W_{P_0,C}}{\int_\Omega W_{P_0,C}(\omega)\, dP_0} = \frac{W_{P_0,C}}{c_{H_{C,1}}}, \tag{14}$$

which is the same for all $x \neq 0$. From (9), we get that the state-independent value of $H_{C,x}$ is $xc_{H_{C,1}}$. QED

Suppose that the loss function L in a statistical decision problem will be paid as $L(\omega, q)$ units of currency C when the agent chooses action q and ω is the state of nature. The agent wants to choose q to maximize

$$\int_\Omega U_P(\omega,\ A_C^{-1}(-L(\omega,\ q)))\, dP(\omega), \tag{15}$$

for some $P \in \mathcal{P}$ (hence, for all $P \in \mathcal{P}$). If $U_P(\omega, \cdot)$ is not linear in its second argument, maximizing expected utility will bear no relationship to minimizing expected loss. For this reason, we would like to deal only with currencies in $\mathcal{C}$. Fortunately, there are many currencies in $\mathcal{C}$. Lemma 4 shows how to construct an element of $\mathcal{C}$ from each pair of positive and negative numeraires.

LEMMA 4. Suppose that there exist both a positive numeraire and a negative numeraire. Then there exist (possibly) other positive and negative numeraires H_+ and H_- and a currency C such that utility is linear in the values of C, both H_+ and H_- have the same corresponding probability/utility pairs, and that common probability/utility pair corresponds to every element of C.

https://doi.org/10.1086/673267 Published online by Cambridge University Press

Proof. Let H'_- be a negative numeraire, and let H'_+ be a positive numeraire. For each probability $P \in \mathcal{P}$, let

$$m_P(\omega) = \min\{U_P(\omega, H'_+(\omega)),\ -U_P(\omega, H'_-(\omega))\}, \tag{16}$$

which is strictly positive for all ω. For each ω, let $z(\omega) = -m_P(\omega)/U_P(\omega, H'_-(\omega))$. It follows from (7) that $z(\cdot)$ is the same for all P and that $0 < z(\omega) \leq 1$. Define $H_-(\omega)$ to be $H'_-(\omega)$ with probability $z(\omega)$ and status quo with probability $1 - z(\omega)$. Similarly, let $w(\omega) = m_P(\omega)/U_P(\omega, H'_+(\omega))$, which is also the same for all P and $0 < w(\omega) \leq 1$. Define $H_+(\omega)$ to be $H'_+(\omega)$ with probability $w(\omega)$ and status quo with probability $1 - w(\omega)$. By construction, we have

$$U_P(\omega, H_+(\omega)) = -U_P(\omega, H_-(\omega)) = m_P(\omega);$$

hence, H_+ and H_- share a common corresponding probability/utility pair as seen from the proof of lemma 1.

For each $-1 \leq x \leq 0$, let $H_x(\omega)$ assign $H_-(\omega)$ with probability $-x$ and status quo with probability $1 + x$. For $0 < x \leq 1$, let $H_x(\omega)$ assign $H_+(\omega)$ with probability x and status quo with probability $1 - x$. Define $C = \{H_x : -1 \leq x \leq 1\}$. First, note that $A_C(H_x) = x$ and $R_C = [-1, 1]$ satisfy definition 4, so that C is a currency. Also, for $-1 \leq x \leq 1$,

$$U_P(\omega, A_C^{-1}(x)) = U_P(\omega, H_x(\omega)) = m_P(\omega)x,$$

for all P and all ω. Since $m_P(\omega) > 0$ for all P and all ω, $W_{P,C}(\omega) = m_P(\omega)$ in definition 5. The final two claims follow from lemma 3 and the facts that $H_- = A_C^{-1}(-1)$ and $H_+ = A_C^{-1}(1)$. QED

The construction in the proof of lemma 4 was first introduced by Smith (1961), who calls it an adaptation from Savage (1954). Intuitively, the method of Smith (1961) is to replace x units of a currency C' with an NM lottery that has probability proportional to $|x|$ of receiving (or paying) a fixed amount and stays in status quo otherwise. The utility of such an NM lottery is proportional to x, regardless of whether $C' \in \mathcal{C}$. In this way, we need only evaluate the utility at a single positive currency value and at a single negative currency value in C'. Next, we show how to make use of this idea in decision problems.

3.3. Paying Loss in a Currency. Consider a statistical decision problem with a set $\mathcal{E}$ of available actions and a bounded loss function $L : \Omega \times E \to \mathbb{R}$. That is, the decision maker's fortune will change to $-L(\omega, q)$ (in some yet to be named currency) if the chosen action is q and the state of nature is ω. Here, q can be a very general action. For example, q can be a function of random variables whose values will not be observed until some

https://doi.org/10.1086/673267 Published online by Cambridge University Press

later time, presumably before the loss gets paid. All that is required is that $L(\omega, q)$ is known in time for paying the loss and that there is enough measurability to be able to compute expected values.

Suppose that we want the loss to be paid using a currency C'. Rather than paying directly in units of C', let H'_- and H'_+ be negative and positive numeraires in C', respectively, and construct the currency C in lemma 4. Define $x(\omega, q) = -L(\omega, q)/M$, where M is an upper bound on the loss function. If the agent chooses action q, change the agent's fortune to $x(\omega, q)$ units of currency C. The agent's expected utility (15) becomes

$$-\frac{1}{M}\int_{\Omega} L(\omega, q) W_{P,C}(\omega)\, dP(\omega). \tag{17}$$

We are now in position to state the following key result.

THEOREM 1. Suppose either that we construct a currency C as in lemma 4 or that utility is already linear in an existing currency C. Suppose also that a decision problem has a loss function $L(\omega, q)$ that is bounded by 1. This means that the agent's fortune moves to $-L(\omega, q)$ units of currency C if the agent chooses action q and the state of nature is ω. Then the agent maximizes expected utility by minimizing expected loss using the probability Q that corresponds to C.

Proof. If the currency C is constructed as in lemma 4, let Q be the probability corresponding to H_-. Then for each $P \in \mathcal{P}$, $dQ/dP(\omega)$ is a positive constant times $W_{P,C}(\omega)$, and (17) is a positive constant times

$$-\int_{\Omega} L(\omega, q)\, dQ(\omega), \tag{18}$$

which is maximized by minimizing expected loss under Q. If utility was already linear in some currency C, then (15) is

$$-\int_{\Omega} L(\omega, q) W_{P,C}(\omega)\, dP(\omega), \tag{19}$$

which is a positive constant multiple of (18). Hence, maximizing expected utility is the same as minimizing expected loss under Q. QED

If we contemplate different choices for the currency in which the loss is paid, the question arises as to whether some currencies are better for a decision problem than others. We turn to that question in section 6.3. In order to

https://doi.org/10.1086/673267 Published online by Cambridge University Press

choose between different currencies, we need a scale on which to compare them. The natural comparison between currencies is their exchange rate, which we consider in section 4.

4. Exchange Rates. An obvious problem with exchange rates, in the presence of nonlinear utility, is as follows. Let C_1 and C_2 be currencies. Even if one unit of C_2 is worth x units of C_1, it does not necessarily follow that two units of C_2 are worth $2x$ units of C_1. Hence, the exchange rate is difficult to define in a manner that matches how it is used in the foreign exchange market, unless utility is linear in both currencies. We begin the discussion of exchange rates by comparing two numeraires and then extend to currencies in which utility is linear.

DEFINITION 6 (Exchange Rates). Let H_1 and H_2 be numeraires. The *conditional exchange rate* from H_1 to H_2 is the function $E_{H_1,H_2}: \Omega \to \mathbb{R}$ equal to the ratio of their state-dependent values, namely,

$$E_{H_1,H_2}(\omega) = \frac{U_P(\omega, H_2(\omega))}{U_P(\omega, H_1(\omega))}, \tag{20}$$

which is the same for all P according to lemma 2. The *marginal exchange rate* from H_1 to H_2 is the ratio of their marginal values $M_{H_1,H_2} = c_{H_2}/c_{H_1}$.

One can think of the marginal exchange rate between two numeraires as their relative values at the present time. In general, when the loss function in a decision problem will be paid at some future time, the relative values of various numeraires might change between now and when the loss is paid. In our discussion of decision problems, we think of the conditional exchange rates between numeraires as their future exchange rates at the time when the loss will be paid.

Notice that $E_{H_2,H_1} = 1/E_{H_1,H_2}$ and $M_{H_2,H_1} = 1/M_{H_1,H_2}$. If H_3 is a third numeraire, then $E_{H_1,H_3} = E_{H_1,H_2}E_{H_2,H_3}$ and $M_{H_1,H_3} = M_{H_1,H_2}M_{H_2,H_3}$, as one would expect of exchange rates. Next, we present some natural relationships between conditional and marginal exchange rates.

LEMMA 5. Let H_1 and H_2 be numeraires with corresponding probability/utility pairs (P_1, U_{P_1}) and (P_2, U_{P_2}). Then

$$E_{H_1,H_2} = \frac{dP_2}{dP_1} M_{H_1,H_2}. \tag{21}$$

Proof. Let (P, U_P) be a probability/utility pair. From the construction in the proof of lemma 1, we see that $dP_i/dP = U_P(\cdot, H_i)/c_{H_i}$ for $i = 1, 2$. It follows that

https://doi.org/10.1086/673267 Published online by Cambridge University Press

$$\frac{dP_2}{dP_1} = \frac{dP_2/dP}{dP_1/dP} = \frac{U_P(\cdot, H_2)c_{H_1}}{U_P(\cdot, H_1)c_{H_2}} = \frac{E_{H_1,H_2}}{M_{H_1,H_2}};$$

hence, (21) holds. QED

LEMMA 6. Under the conditions of lemma 5, $M_{H_1,H_2} = P_1(E_{H_1,H_2})$.

Proof. From lemma 5,

$$P_1(E_{H_1,H_2}) = P_1\left(\frac{dP_2}{dP_1}M_{H_1,H_2}\right) = M_{H_1,H_2}. \tag{22}$$

QED

In words, lemma 6 says that the marginal exchange rate from H_1 to H_2 is the mean of the conditional exchange rate with respect to the probability corresponding to the utility that gives H_1 state-independent values.

The remaining results in this section concern the collection $\mathcal{C}$ of currencies such that utility is linear in each of the currencies. In the notation of lemma 3, the conditional exchange rate between x units of two different currencies C_1 and C_2 in $\mathcal{C}$ is

$$E_{H_{C_1,x},H_{C_2,x}}(\omega) = \frac{U_P(\omega, H_{C_2,x}(\omega))}{U_P(\omega, H_{C_1,x}(\omega))} = \frac{W_{P,C_2}(\omega)}{W_{P,C_1}(\omega)}, \tag{23}$$

for all $x \neq 0$. That is, as long as we compare numeraires consisting of the same numerical amounts x of currency, the conditional exchange rate does not depend on the common amount x. Lemma 2 shows that $E_{H_{C_1,x},H_{C_2,x}}$ does not depend on P, which fact also follows quickly from (11). Use the symbol $E_{C_1,C_2}(\omega)$ to denote the conditional exchange rate in (23). Let $M_{C_1,C_2} = c_{H_{C_2,1}}/c_{H_{C_1,1}}$ stand for the marginal exchange rate from C_1 to C_2.

In the linear case, exchange rates have interpretations much like what we see in foreign exchange. The marginal exchange rate M_{C_1,C_2} is the number of units of C_1 that has the same value as one unit of C_2 at the present time. The conditional exchange rate has the same interpretation state by state.

We are now in a position to see how the Bayesian in examples 1 and 2 can clear up the inconsistent choices that were made.

EXAMPLE 9. Reconsider examples 1 and 2. We have not yet given enough information to determine the probabilities that correspond to each of the two currencies. But we know that they are not the same. First suppose that $P_1(\{1\}) = P_1(\{2\}) = 0.5$ is the probability that corresponds to C_1. Let P_2 be the probability that corresponds to C_2. According to lemma 5,

$$\frac{dP_2}{dP_1} = \frac{E_{C_1,C_2}}{M_{C_1,C_2}}.$$

https://doi.org/10.1086/673267 Published online by Cambridge University Press

In example 2, we specified $E_{C_1,C_2}(1) = 0.5$ and $E_{C_1,C_2}(2) = 2$. This makes $M_{C_1,C_2} = 0.5 \times 0.5 + 0.5 \times 2 = 1.25$, and

$$\frac{dP_2}{dP_1}(\omega) = \frac{1}{1.25} \times \begin{cases} 0.5 & \text{if } \omega = 1, \\ 2 & \text{if } \omega = 2. \end{cases} = \begin{cases} 0.4 & \text{if } \omega = 1, \\ 1.6 & \text{if } \omega = 2. \end{cases}$$

So $P_2(\{1\}) = 0.5 \times 0.4 = 0.2$ and $P_2(\{2\}) = 0.5 \times 1.6 = 0.8$.

Using currency C_1, the expected losses for actions a and b are as given in example 1, namely, 1.25 and 1 respectively, and the agent chooses b. Using currency C_2, the expected loss for action a is again 1, while the expected loss for action b is $2 \times 0.2 + 0.5 \times 0.8 = 0.8$, and the agent still chooses b, as expected.

For completeness, suppose next that the probability corresponding to C_2 is $Q_2(\{1\}) = Q_2(\{2\}) = 0.5$, which happens to be the same as P_1 above. Let Q_1 be the probability corresponding to C_1. The conditional exchange rate that we need now is $E_{C_2,C_1} = 1/E_{C_1,C_2}$; that is, $E_{C_2,C_1}(1) = 2$, $E_{C_2,C_1}(2) = 0.5$. The marginal exchange rate is now $M_{C_2,C_1} = 0.5 \times 2 + 0.5 \times 0.5 = 1.25$, and

$$\frac{dQ_1}{dQ_2} = \frac{E_{C_2,C_1}}{M_{C_2,C_1}} = \begin{cases} 1.6 & \text{if } \omega = 1, \\ 0.4 & \text{if } \omega = 2. \end{cases}$$

So, $Q_1(\{1\}) = 0.5 \times 1.6 = 0.8$ and $Q_1(\{2\}) = 0.5 \times 0.4 = 0.2$. Using currency C_1, the expected losses for actions a and b are $0.5 \times 0.8 + 2 \times 0.2 = 0.8$ and 1, respectively. The agent chooses a. Using currency C_2, the expected losses are 1 and 1.25 (as in example 2), and the agent chooses a again, as expected. We do not know how the minimax decision maker can resolve the inconsistent choices in examples 1 and 2, and we leave it as an open question.

5. When Currency Does Not Matter. There are cases in which the currency used for charging a loss does not affect the decision.

LEMMA 7. Assume the conditions of theorem 1. Let C_1 and C_2 be two currencies in $\mathcal{C}$ with corresponding probability/utility pairs (P_1, U_{P_1}) and (P_2, U_{P_2}) and with a marginal exchange rate equal to 1. Then, the following statements are equivalent:

- for each q, the two expected utilities from paying the loss in units of C_1 and C_2 are equal,
- $L(\cdot, q)$ is uncorrelated with E_{C_1,C_2} under P_1 for all q,
- $L(\cdot, q)$ is uncorrelated with E_{C_2,C_1} under P_2 for all q.

https://doi.org/10.1086/673267 Published online by Cambridge University Press

Proof. First, we show that the second bullet implies the first and third bullets. Suppose that $L(\cdot, q)$ is uncorrelated with E_{C_1,C_2} under P_1 for all q. According to lemma 5, $dP_2/dP_1 = E_{C_1,C_2}M_{C_2,C_1}$, so $L(\cdot, q)$ is uncorrelated with dP_2/dP_1 under P_1. Then, for each q,

$$P_2(L(\cdot, q)) = \int_\Omega L(\omega, q)dP_2(\omega) = \int_\Omega L(\omega, q)\frac{dP_2}{dP_1}(\omega)\,dP_1(\omega)$$
$$= P_1(L(\cdot, q))P_1\left(\frac{dP_2}{dP_1}\right) = P_1(L(\cdot, q)),$$

where the third equality follows from $L(\cdot, q)$ and dP_2/dP_1 being uncorrelated under P_1. This establishes the first bullet.

Next, we show that $L(\cdot, q)$ is uncorrelated with E_{C_2,C_1} under P_2, which is equivalent to showing that $L(\cdot, q)$ is uncorrelated with dP_1/dP_2 under P_2. We have

$$P_2\left(L(\cdot, q)\frac{dP_1}{dP_2}\right) = P_1(L(\cdot, q)) = P_2(L(\cdot, q)) = P_2(L(\cdot, q))P_2\left(\frac{dP_1}{dP_2}\right),$$

which is the third bullet. That the third bullet implies the first two follows by repeating the above argument with subscripts 1 and 2 switched.

To complete the proof, it suffices to show that the first bullet implies the second bullet. Suppose that $P_2(L(\cdot, q)) = P_1(L(\cdot, q))$ for all q. Since E_{C_1,C_2} is a constant times dP_2/dP_1, we need to show that

$$P_1\left(L(\cdot, q)\frac{dP_2}{dP_1}\right) = P_1(L(\cdot, q))P_1\left(\frac{dP_2}{dP_1}\right). \tag{24}$$

The left side of (24) is $P_2(L(\cdot, q))$, and the right side is $P_1(L(\cdot, q))$, which are equal. QED

The ability to apply lemma 7 depends on how complicated the loss function is and how complicated the decision rules $q \in \mathcal{E}$ are. If all of the random variables that go into determining the loss (and q) are independent of E_{C_1,C_2} under P_1, then $L(\cdot, q)$ is uncorrelated with E_{C_1,C_2} under P_1 for all q, and the lemma says that all actions will be ranked the same, regardless of which currency (C_1 or C_2) is used to pay the loss. Put less technically, if the decision problem is independent of the exchange rate, then it does not matter what currency is used for charging the loss.

6. Elicitation via Proper Scoring Rules

6.1. Elicitation as a Decision Problem. Proper scoring rules were designed to give experts the proper incentives for providing their subjective

https://doi.org/10.1086/673267 Published online by Cambridge University Press

probabilities and expected values when being elicited. Being scored by a proper scoring rule is a special case of a statistical decision problem.

> DEFINITION 7 (Proper Scoring Rule). Let R be a set of real numbers and let $(\mathcal{X}, D)$ be a measurable space. Let $g: \mathcal{X} \times R \to [0, \infty]$ be a function such that $g(x, q)$ is measurable in x for all q. For each probability Q over Ω and each bounded random variable X, let $Q(X)$ denote the mean of X. Suppose that, for every Q and every X, $Q[g(X, q)]$ is minimized as a function of q at $q = Q(X)$. Then g is a *proper scoring rule*. If, for every Q, $q = Q(X)$ is the unique minimizer, then g is *strictly proper*.

Definition 7 could be extended to allow unbounded random variables, but then one has to deal with the possibility of infinite or undefined means. Although the definition of proper scoring rule above is widely accepted, there is some controversy about what should be called strictly proper. Some authors reserve the qualification "strictly proper" for scoring rules such that are designed to elicit an entire distribution. That is, R in definition 7 is a set of probability measures, and g is strictly proper if and only if $q = Q$ is the unique minimizer of $Q[g(X, q)]$ for every $Q \in R$. Gneiting (2011a) calls a scoring rule *strictly consistent* if it satisfies the final clause of definition 7. For the remainder of this article, we will continue to follow definition 7, which matches the usage in Gneiting (2011b). The two definitions agree when R is a set of Bernoulli distributions.

Suppose that we wish to learn a particular agent's subjective expectation for a random variable X (possibly the indicator of an event). Let g be a strictly proper scoring rule. We can create a statistical decision problem with loss function $L(\omega, q) = g(X(\omega), q)$. If we were able to convince the agent to provide us with the value q that minimizes $\int_\Omega g(X(\omega), q)\, dP(\omega)$, where P is the agent's subjective probability distribution, we would learn $P(X)$, according to the definition of proper scoring rule. But theorem 1 says that the solution to a statistical decision problem depends on which currency is used for charging the loss (score). If the agent is being given the proper incentive for providing his or her subjective probability of an event, then how can the elicited probability depend on which currency is being used for scoring?

The answer is straightforward. Minimizing expected score is the same as maximizing expected utility, only when the probability/utility pair corresponds to the currency used for scoring. Using different currencies creates different decision problems for the agent. These different decision problems can have different solutions. The confusion arises from mistakenly treating the scores as pure (unit-less) numbers, regardless of how their values change from state to state.

Chapter 6 of Degroot (1970) constructs a unique subjective probability P without reference to changes in an agent's fortunes. Then chapter 7 of

https://doi.org/10.1086/673267 Published online by Cambridge University Press

Degroot (1970) goes on to construct a state-independent utility. A more complicated theory of Karni, Schmeidler, and Vind (1983) and Schervish, Kadane, and Seidenfeld (1991) allows the construction of a unique probability and state-dependent utility (P, U_P). Suppose that an agent believes that he or she has constructed a unique probability and utility (P, U_P) by one of these methods or by some other method. If we attempt to elicit $P(X)$ from this agent using a strictly proper scoring rule, we will not be successful, unless we already know enough about U_P to choose a currency with state-independent values.

Lemma 7 has a simpler form when restricted to elicitation via proper scoring rules.

> PROPOSITION 2. Let C_1 and C_2 be two currencies with corresponding probability/utility pairs (P_1, U_{P_1}) and (P_2, U_{P_2}). Then the following items are equivalent:
>
> - $P_1(X) = P_2(X)$,
> - X is uncorrelated with E_{C_1,C_2} under P_1,
> - X is uncorrelated with E_{C_2,C_1} under P_2.

The proof of proposition 2 is similar to that of lemma 7 and will not be given.

One can elicit other aspects of a probability distribution, such as quantiles, using other loss functions. For example, if $P(|X|) < \infty$, then $\int_\Omega |X(\omega) - q| dP(\omega)$ is minimized over q by setting q equal to any median of the distribution of X under P. For more general quantiles, one can use loss functions of the form

$$L(\omega, q) = \begin{cases} a[X(\omega) - q] & \text{if } X(\omega) > q, \\ b[q - X(\omega)] & \text{if } X(\omega) \le q, \end{cases} \tag{25}$$

where $a, b > 0$. In this case, $\int_\Omega L(\omega, q)\, dP(\omega)$ is minimized by setting q equal to any $a/(a + b)$ quantile of the distribution of X under P. In such a decision problem, if the loss is settled in a currency, the quantile elicited will be from the probability corresponding to the currency.

6.2. Strategic Choices in Elicitation. De Finetti (1974, 93) noticed a shortcoming of the use of gambles for elicitation and preferred to use a proper scoring rule. When an agent is gambling, there is an opponent who gets to choose which side of the gamble to take. Suppose that the agent has reason to believe that the opponent has a higher mean for the random variable of interest than does the agent. Then the agent will have an in-

https://doi.org/10.1086/673267 Published online by Cambridge University Press

centive to specify a slightly higher value than his or her true mean. Here is an example.

> EXAMPLE 10. In the gambling formulation of elicitation, the agent is asked to specify the mean μ of a random variable X, with the understanding that the agent then feels that it is fair to receive $\alpha[X(\omega) - \mu]$ in state ω, where the real scalar α is chosen by an opponent. For example, suppose that the agent thinks that $\mu = 0.6$ meets the above condition, but the agent is certain that the opponent would choose $\mu \geq 0.8$ if it were up to the opponent. So, the agent feels that $[X(\omega) - 0.7]$ would be advantageous to receive, as it is 0.1 higher than the fair value of $[X(\omega) - 0.6]$. Also, the opponent (in the opinion of the agent) would think it is advantageous to receive $-[X(\omega) - 0.7]$, as it is 0.1 higher than the fair $-[X(\omega) - 0.8]$. In this case, the agent has an incentive to specify a value of μ that is higher than his or her mean of X; hence, the gambling formulation would fail to provide a proper elicitation in this case.

Strategic considerations of the sort in example 10 can undermine the value of gambling as an elicitation method. Scoring rules do not involve an opponent who has any decisions that should influence the agent. For this reason, de Finetti believed that scoring rules should be preferred to gambles as a means of elicitation. Using gambles to elicit probabilities in finite spaces also leads to the same dependence on currency that we noticed above, as shown by Schervish, Seidenfeld, and Kadane (1990). In this article, we have extended the results of Schervish et al. (1990) to proper scoring rules as well as to all statistical decision problems, even in general spaces. Next, we show that a different set of strategic considerations arises when using scoring rules to elicit probabilities.

6.3. Strategic Choice of Currency. Suppose that we wish to elicit the mean of a random variable from an agent who is given the choice of currency in which to be scored before announcing the mean. Are some currency choices better than others? Without further conditions, the answer is an obvious yes. Surely it is better to pay a score of x units in pennies than to pay x units in dollars. To avoid such trivial answers, we need to standardize currencies somehow and compare only those currencies that are of the same size according to the standardization. But even that appears not to be enough to prevent strategic choice of currency.

> EXAMPLE 11. Let $\Omega = (0, 1)$. Suppose that we are trying to elicit the probability of the event $F = (0, 1/2)$. That is, $X(\omega)$ is the indicator of F. Suppose that we are using Brier score, $g(x, q) = (x - q)^2$. Suppose also that we assume that utility is linear in all of the currencies that we use in

https://doi.org/10.1086/673267 Published online by Cambridge University Press

this example. Let P_1 be the probability corresponding to a currency C_1 having state-independent values, and let $W_{P_1,C_1} = 1$. Suppose that P_1 is the uniform distribution on $(0, 1)$. If the agent chooses to be scored in currency C_1, then $q = P_1(F) = 1/2$, and the expected Brier score is $\mathrm{Var}_{P_1}(X) = 1/4$.

Each alternative currency C_2 corresponds to a conditional exchange rate $E_{C_1,C_2}(\cdot) = W_{P_1,C_2}(\cdot)$ that is integrable with respect to P_1. The corresponding probability P_2 that gives currency C_2 state-independent values has $dP_2/dP_1 = W_{P_1,C_2}/c$, where $c = P_1[W_{P_1,C_2}] = M_{C_1,C_2}$ is the marginal exchange rate. If the agent chooses to be scored in currency C_2, the probability is

$$P_2(F) = \int_0^{1/2} \frac{W_{P_1,C_2}(\omega)}{c}\, d\omega.$$

For example, suppose that we consider a currency C_2 with $W_{P_1,C_2}(\omega) = 2\omega = dP_2/dP_1$ so that $c = 1$. Then $P_2(F) = 1/4 = q$, and the expected Brier score under P_2 is $\mathrm{Var}_{P_2}(X) = 3/16$. Since the marginal exchange rate is 1, paying 3/16 units of C_2 is preferred to paying 1/2 unit of C_1.

Taking the above comparison between C_1 and C_2 further, let C_n be a currency with $W_{P_1,C_n}(\omega) = n\omega^{n-1}$. Then $P_n(F) = 1/2^n$, $M_{C_1,C_n} = 1$, and $\mathrm{Var}_{P_n}(X) = (1 - 2^{-n})/2^n$. The differences between the expected scores in currencies C_1 and C_n cannot be explained by the marginal exchange rate between the two currencies since the marginal exchange rates are all 1. What is happening is that C_n is essentially worthless (when measured in units of C_1) if F occurs. The agent announces a very small probability of F and agrees to pay a large score in currency C_n if F occurs. But such a large score is not worth much in other currencies. If F^C occurs, making C_n more valuable relative to other currencies, the agent does not have to pay very much in units of C_n because $P_n(F^C)$ is close to 1.

The same strategic consideration, that is, choice of currency, does not arise when using gambles for elicitation. In that formulation, the gambles to which an agent commits are all fair, regardless of in what currency they are settled. Without a secondary criterion with which to distinguish fair gambles, there is no way to choose among them.

6.4. Choice of Scoring Rule. Another strategic consideration arises if the agent is given the choice of which scoring rule will be used to score the elicitation. Clearly, scaling a scoring rule down is advantageous to the agent being scored. In order to compare scoring rules that are comparable in terms of the payout, we need an appropriate standardization. One naive standardization is to scale by $\sup_{x,q} g(x, q) = 1$.

https://doi.org/10.1086/673267 Published online by Cambridge University Press

For simplicity, consider the case in which X is the indicator of some event. According to theorem 4.2 of Schervish (1989), every bounded left-continuous strictly proper scoring rule with $g(x, x) = 0$ for $x \in \{0, 1\}$ has the form

$$g(x,q) = \begin{cases} \int_{[0,q)} p\, d\lambda(p) & \text{if } x = 0, \\ \int_{[q,1)} (1-p)\, d\lambda(p) & \text{if } x = 1, \end{cases} \tag{26}$$

for some measure λ on $[0, 1]$ that assigns a positive measure to every nondegenerate interval. In order for $\sup_{x,q} g(x, q) = 1$, we need λ to be two times a probability that has mean of $1/2$. By choosing λ to put as much of its mass as possible near the two extreme values of 0 and 1, the expected score can be made as close as one likes to 0, no matter what q happens to be. Hence, the agent would like to be scored by a rule corresponding to such a λ, regardless of the currency.

An alternative normalization of scoring rules is to use the maximin expected score. That is, normalize by $\sup_q [qg(1, q) + (1 - q)g(0, q)]$. In this case, the expected score will lie on a strictly concave curve m on $[0, 1]$, with a maximum value of 1 and satisfying $m(0) = m(1) = 0$. If $m(q_0) = 1$, then the curve m lies strictly above the piecewise linear function $f(q) = \min\{(1 - q_0)q, (1 - q)q_0\}$ (except for $q \in \{0, q_0, 1\}$, where $f(q) = m(q)$). We can make $m(q)$ arbitrarily close to $f(q)$, by making λ concentrate its mass arbitrarily close to q_0. In such a case, if the agent's subjective probability of the event being forecast is q, then the best expected score will be approximately $f(q)$, which will be minimized by choosing λ so that $q_0 = 1$ if $q < 1/2$ and $q_0 = 0$ if $q > 1/2$. If $q = 1/2$, either $q_0 = 0$ or $q_0 = 1$ will do equally well. If the agent also gets to choose the currency along with the scoring rule, he or she would choose a currency such that q is as close to 0 or 1 as is feasible and match it with a scoring rule that made the optimal expected score as close as possible to 0 near that q. In the gambling framework, there is no obvious strategic counterpart to the choice of the scoring rule on the part of the agent being scored.

6.5. Converting between Currencies. Our results show that a mean elicited by a scoring rule comes from the probability P associated with the utility U_P that gives state-independent values to the currency used for eliciting. In general, it is not possible to infer $Q(X)$ from $P(X)$, even if we know the conditional exchange rate between the two currencies C_P and C_Q

https://doi.org/10.1086/673267 Published online by Cambridge University Press

that have state-independent values under P and Q, respectively. Even when X is the indicator of an event F, we have

$$Q(F) = \int_F \frac{dQ}{dP}(\omega)\, dP(\omega).$$

It is true that

$$\frac{dQ}{dP} = \frac{E_{C_P,C_Q}}{M_{C_P,C_Q}}, \tag{27}$$

but we still need to know P for all subsets of the event F, not just $P(F)$ (unless E_{C_P,C_Q} is constant over F or F^C). In general,

$$Q(X) = P\left(X\frac{dQ}{dP}\right) = \frac{P(XE_{C_P,C_Q})}{M_{C_P,C_Q}}, \tag{28}$$

which can also be written as $P(XE_{C_P,C_Q}) = Q(X)P(E_{C_P,C_Q})$.

6.6. Finite State Spaces. In finite state spaces, we can make a bit more progress. Let $\Omega = \{\Omega_1, \ldots, \Omega_n\}$. Schervish et al. (1990) dealt with this case, and the events whose probabilities were being elicited (via gambles) were singletons $\{\omega_1\}, \ldots, \{\omega_n\}$. In such cases, E_{C_P,C_Q} is constant on singletons; hence, we can convert probabilities of singletons from one currency to the next if both the conditional exchange rate and the marginal exchange rate are known. If the probabilities of all singletons are elicited in the same currency, then the marginal exchange rate can be computed from the conditional exchange rate. If the probability of each singleton is elicited in a (possibly) different currency, one can set up a system of equations whose solution will give the necessary marginal exchange rates.

To be specific, suppose that the probability of $\{\omega_i\}$ is elicited in currency C_i for $i = 1, \ldots, n$ with corresponding probabilities $P_1, \ldots, P_n$. Let C_0 be a currency with corresponding probability P_0. We assume that we know E_{C_0,C_i} for all i, even if we do not know M_{C_0,C_i}. So, we elicit $P_i(\{\omega_i\}) = p_i$ for $i = 1, \ldots, n$. As in (27),

$$P_0(\{\omega_i\}) = P_i(\{\omega_i\})\frac{E_{C_i,C_0}(\omega_i)}{M_{C_i,C_0}} = p_i\frac{M_{C_0,C_i}}{E_{C_0,C_i}(\omega_i)}. \tag{29}$$

For each i, we can set up an equation giving the value of M_{C_0,C_i}. According to (22) and then (29),

$$M_{C_0,C_i} = \sum_{j=1}^{n} P_0(\{\omega_j\})E_{C_0,C_j}(\omega_j) = \sum_{j=1}^{n} p_j M_{C_0,C_j}. \tag{30}$$

https://doi.org/10.1086/673267 Published online by Cambridge University Press

For $i = 1, \ldots, n$, (30) gives us n linear equations in (at most) n unknowns M_{C_0,C_i} for $i = 1, \ldots, n$. The equations are linearly dependent, and every scalar multiple of each solution is also a solution. The appropriate scaling can be determined from the fact that $\sum_{i=1}^{n} P_0(\{\omega_i\}) = 1$. If two or more events were elicited in the same currency, there will be further linear dependence, which could be removed by using only one equation for each unique currency.

7. Discussion. This article explores implications of the fact that preferences between Anscombe and Aumann (1963) style horse lotteries cannot reveal a unique probability and state-independent utility in the presence of numeraires with state-dependent values. The most that one can determine is a state-dependent expected utility representation in the form of definition 1. Even an agent who believes that they have a unique probability and state-dependent utility (P, U_P) cannot ignore the units in which losses are paid in a statistical decision problem. The formal solution to a decision problem as if the losses were unit-less pure numbers will not match the agent's solution, unless the loss is paid in a currency that has state-independent values under U_P.

Schervish et al. (1990) considered elicitation via gambles, but the current article shows that general statistical decision problems suffer from possible state dependence of the currency used for charging the loss. When one pays the loss in a particular currency C, then a Bayesian will solve the decision problem using the probability Q, where (Q, U_Q) is the particular state-dependent utility representation of the agent's preferences such that currency C has state-independent values under U_Q. If one changes the currency to C' and $(Q', U_{Q'})$ is an equivalent state-dependent utility for which C' has state-independent values under $U_{Q'}$, then the agent will solve the decision problem using probability Q'. Lemma 7 and proposition 2 give conditions under which the solutions to various decision problems will not depend on the currency in which the loss is paid.

We examine elicitation of subjective probability in detail because the implications of state-dependent utility are so striking for elicitation. Under some restrictive conditions, one can convert means elicited in one currency to means that would have been elicited in another currency. But, one cannot elicit an agent's "true subjective probability" (whatever that means) unless that probability happens to correspond to the state-dependent utility that gives state-independent values to the currency in which one does the elicitation. If the random variable whose prevision is being elicited is uncorrelated with the conditional exchange rate between two currencies, then the same prevision will be elicited using either currency. If one is to take seriously the idea that elicited probabilities stand for something that can be used for statistical inference, one needs to be confident that those probabil-

https://doi.org/10.1086/673267 Published online by Cambridge University Press

ities were derived in a manner consistent with their intended use. If it is possible to constrain the effects of the decisions so that they do not involve fortunes whose relative values vary from state to state, then one can feel safe that probabilities elicited using such fortunes as currency values will be meaningful. The challenge is making sure that the decision problem is so constrained.

Additional work is needed in order to identify the implications of state dependence on decision problems with multiple decision points. For example, an agent may get the opportunity to revise a decision after learning additional information. The agent may be asked to make several decisions at different times. It is well known that exchange rates change over time, and it makes sense to model exchange rates as stochastic processes. In this article, we have considered only two times, namely, when the decision is made and when the loss is paid. We also believe that state dependence has implications for financial product pricing, especially in the foreign exchange market.

Finally, we have presented some results concerning strategic choices that an agent might make when being scored, but we have not yet studied strategic choices that are available to the elicitor who is requesting the elicited prevision. Such a study could proceed if we specified how the elicited prevision was going to affect the elicitor as well as his or her state-dependent utility representation and his or her opinion of the agent. This problem is left for future study.

REFERENCES

Anscombe, F., and R. Aumann. 1963. "A Definition of Subjective Probability." *Annals of Mathematical Statistics* 34:99–205.

Braithwaite, R. B. 1955. *Theory of Games as a Tool for the Moral Philosopher*. Cambridge: Cambridge University Press.

De Finetti, B. 1974. *Theory of Probability*. 2 vols. New York: Wiley.

Degroot, M. 1970. *Optimal Statistical Decisions*. New York: Wiley.

Fisburn, P. 1970. *Utility Theory for Decision Making*. New York: Wiley.

Gneiting, T. 2011a. "Making and Evaluating Point Forecasts." *Journal of the American Statistical Association* 106:746–62.

———. 2011b. "Quantiles as Optimal Point Forecasts." *International Journal of Forecasting* 27:197–207.

Harsanyi, J. C. 1977. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge: Cambridge University Press.

Karni, E., D. Schmeidler, and K. Vind. 1983. "On State Dependent Preferences and Subjective Probabilities." *Econometrica* 51:1021–31.

Rubin, H. 1987. "A Weak System of Axioms for 'Rational' Behavior and the Non-separability of Utility from Prior." *Statistics and Decisions* 5:47–58.

Savage, L. 1954. *Foundations of Statistics*. New York: Wiley.

———. 1971. "Elicitation of Personal Probabilities and Expectations." *Journal of the American Statistical Association* 66:783–801.

https://doi.org/10.1086/673267 Published online by Cambridge University Press

Schervish, M. J. 1989. "A General Method for Comparing Probability Assessors." *Annals of Statistics* 17:1856–79.
Schervish, M. J., J. B. Kadane, and T. Seidenfeld. 1991. "Shared Preferences and State-Dependent Utilities." *Management Science* 37:1575–89.
Schervish, M. J., T. Seidenfeld, and J. B. Kadane. 1990. "State-Dependent Utilities." *Journal of the American Statistical Association* 85:840–47.
Seidenfeld, T., M. J. Schervish, and J. B. Kadane. 2009. "Preference for Equivalent Random Variables: A Price for Unbounded Utilities." *Journal of Mathematical Economics* 45:329–40.
Shimony, A. 1955. "Coherence and the Axioms of Confirmation." *Journal of Symbolic Logic* 20:1–28.
Smith, C. A. B. 1961. "Consistency in Statistical Inference and Decision." *Journal of the Royal Statistical Society* B 23 (1): 1–37.
Von Neumann, J., and O. Morgenstern. 1947. *Theory of Games and Economic Behavior*. 2nd ed. Princeton, NJ: Princeton University Press.
Wald, A. 1950. *Statistical Decision Functions*. New York: Wiley.

https://doi.org/10.1086/673267 Published online by Cambridge University Press