

Coherent Decision Analysis With Inseparable Probabilities and Utilities

ROBERT F. NAU

The Fuqua School of Business, Duke University, Durham, North Carolina 27708-0120

Abstract

This article explores the extent to which a decision maker's probabilities can be measured separately from his/her utilities by observing his/her acceptance of small monetary gambles. Only a partial separation is achieved: the acceptable gambles are partitioned into a set of "belief gambles," which reveals probabilities distorted by marginal utilities for money, and a set of "preference gambles," which reveals utilities reciprocally distorted by marginal utilities for money. However, the information in these gambles still enables us to solve the decision maker's problem: his/her utility-maximizing decision is the one that avoids arbitrage (i.e., incoherence or Dutch books).

Key words: coherence, subjective probability, state-dependent utility, small worlds, risk neutral probabilities, noncooperative games, arbitrage, Dutch books

1. Introduction

One of the cornerstones of the theory of rational choice is the separation of *beliefs* from *preferences*. Beliefs are based primarily on information, so the story goes, and are represented by numerical probabilities attached to events. Preferences are based on personal tastes and are represented by numerical utilities attached to consequences of events and decisions. Thus equipped with probability distributions to summarize their information and utility functions to summarize their tastes, rational economic persons set out to maximize their expected utility. This representation of rational behavior is taken very literally by professional decision analysts, who attempt to elicit probabilities and utilities from their clients in order to clarify their decisions, and also by economic theorists, who model behavior in games and markets by making explicit assumptions about the probabilities and utilities of the agents in them. For example, game theoretic models usually assume that the players' utilities are common knowledge and that their probability distributions are mutually consistent. Rational-expectations models of market phenomena assume that the probabilities of different agents are not only mutually consistent but are "correct" in the sense of the agreeing with the predictions of the economist's model.

Given that holding the correct probabilities and knowing the values of other persons' utilities is so important in theory, it is remarkable that these quantities are so rarely

displayed in practice. Personal decisions and economic transactions are normally evaluated in terms of money rather than roulette-wheel lotteries or utiles. Money provides not only a medium of exchange but also the language of *prices* through which we articulate our beliefs and preferences to ourselves and to others. We do not routinely assess the relative utilities of Fords, Toyotas, and BMWs, yet we can determine how much more we would pay for one than another. The financial pages do not report a probability distribution for the future value of IBM stock, but they do list the market-clearing prices for options on that stock at different exercise values. In fact, given the importance of money in everyday economic decisions, it is perhaps remarkable that money plays no distinguished role in the foundations of decision theory and microeconomics.

The low profile of probabilities and utilities in everyday economic discourse may reflect an intrinsic difficulty in separating the effects of one from the effects of the other when observing materially significant¹ choices that agents make under conditions of uncertainty. Such choices reveal only the relative product of probability and utility: expected utility equals probability times utility, summed over states of nature. However, an agent's utilities for consequences that we can measure (e.g., money, commodities, services) may vary according to the state of nature in which they are obtained, either because the agent has significant but unobservable prior stakes in the outcomes of events, or because his/her utilities are intrinsically state-dependent.² Thus, for example, two agents may hold very different probability distributions for the future value of a stock price and yet be willing to pay the same price for an option on that stock, because their differences in probabilities are offset by differences in state-dependent utilities for wealth due to prior stakes in the company (stock ownership, employment, etc.). Unless we know a priori how each agent's utilities vary across states, we cannot uniquely separate their probabilities from their utilities by observing their choices under uncertainty. Schervish, Seidenfeld, and Kadane (1990) point out that this inseparability casts doubt on much of probability theory and statistical theory, because it "deals solely with probabilities and not with utilities." The obvious corollary is that it also casts doubt on decision analysis and economic models which require unique determination of an agent's utilities.

This article starts from the premise that an agent's probabilities may be inseparable from his/her utilities and explores the implications for decision analysis. (The implications for economic models—especially noncooperative games—are considered more fully in Nau, 1992b.) We assume that in some "grand world" of events, decisions, and their tangible or intangible consequences, an agent wishes to behave in accordance with the standard axioms of rational choice—i.e., he/she wishes to maximize expected utility with respect to some probability distribution and utility function. We then suppose that the agent's behavior can be observed only in the "small world" of low-stakes monetary side-gambles, and we ask whether these observations are sufficient to determine the optimal grand-world decision. That is, we attempt to carry out a decision analysis by asking the agent to articulate the set of small gambles he/she would be willing to accept in addition to the grand-world consequences he/she faces.

Under these conditions, we find that only a partial separation between probabilities and utilities is achieved: the set of all acceptable gambles can be partitioned into a set of "belief gambles" and a set of "preference gambles." Belief gambles reflect probabilities distorted by state-dependent marginal utilities for money, and are those gambles which

remain acceptable even if the decision is chosen by an arbitrary mechanism. Preference gambles reflect grand-world utilities reciprocally distorted by state-dependent marginal utilities for money, and are those gambles which remain acceptable even if the decision maker receives arbitrary new information before making a decision. Despite the fact that they do not reveal “true” probabilities and utilities, these two types of gambles can be combined in such a way as to reveal the utility-maximizing decision in the grand-world problem. This finding shows that the divide-and-conquer approach of decision analysis is still possible even when probabilities and utilities are inseparable: one merely needs to elicit distorted utilities to pair up with distorted probabilities. However, it also suggests that distorted, rather than true, probabilities and utilities are likely to be the objects of consensus or common knowledge in economic systems.

We also show that expected-utility-maximizing behavior in the grand-world problem corresponds to *coherence* (Dutch book avoidance or “no arbitrage”) in the small-world problem. This result is an extension of de Finetti’s (1937, 1974) formulation of subjective probability theory. Recognizing that prices are the most familiar medium for quantifying subjective judgments, de Finetti defined personal probabilities of events as prices of lottery tickets—i.e., monetary gambling odds—and then showed that such prices must conform to the laws of probability if they are not to lead to a Dutch book. For example, if $P(A)$ is the price of a ticket that pays \$1 if A occurs and \$0 otherwise, $P(B)$ is the price of a similar ticket on a mutually exclusive event B , and $P(A \cup B)$ is the price of a ticket that pays \$1 if *either* event occurs, then the prices must satisfy $P(A \cup B) = P(A) + P(B)$ or else arbitrage is possible. This additive property is a consequence of the duality theorem of linear programming, which underlies many other results in economic theory relating the existence of a price system to the absence of arbitrage (Ross, 1976; Varian, 1987; Nau and McCardle, 1991). Though elegant and practical, the monetary-gamble definition of probability has two important drawbacks. First, it assumes that lottery ticket prices are valid for arbitrary quantities bought or sold, which entails “rigidity in the fact of risk”—i.e., linear utility for money within states. Second, it ignores the possibility that utilities for money may vary between states due to prior stakes or other asymmetries.

In this article, we admit the possibility of nonlinear utility for money within states by introducing the concept of “ ϵ -acceptability,” i.e., the acceptability of gambles for infinitesimal stakes. Further, we accommodate the distortion of probabilities by state-dependent utilities for money by gathering complementary information about utilities through the novel device of preference gambles. The Dutch book argument is then applied to ϵ -acceptable belief and preferences gambles to show that they must conform to the laws of subjective expected utility. These innovations address the objection that often has been levelled against the Dutch book argument as a basis for rational choice theory, namely, that it is incompatible with risk-averse behavior (e.g., Yaari, 1985). We conclude that Dutch book avoidance is precisely what obtains when subjective-expected-utility-maximizing behavior in the grand world is projected into the small world of low-stakes monetary gambles.

The remainder of the article is organized as follows. Section 2 describes the connection between subjective expected-utility maximization and the ϵ -acceptance of monetary gambles. Section 3 points out the intrinsic difficulty in separating probabilities from utilities via material measurements, but shows that a partial separation can be achieved

via the devices of belief and preference gambles. Section 4 describes how, in principle, the parameters of belief and preference gambles can be elicited, and gives consistency conditions which these parameters must satisfy. Section 5 presents the basic duality theorem showing that expected-utility maximization in the “grand world” corresponds precisely to coherence (no arbitrage or Dutch books) in the “small world” of monetary gambles. This reconciles the Dutch book argument with nonlinear utility for money and yields an alternative method for determining an optimal decision. Section 6 illustrates how the preceding elicitation and optimization procedures can be used to solve a simplified decision analysis problem. Section 7 discusses some of the implications of these results for statistical and economic models.

2. Gambling and subjective expected-utility maximization

Consider a decision maker (DM) who faces a set D of available decisions (acts) and a set S of states of nature about whose resolution she is uncertain. For the purposes of this article, these sets will be assumed to be finite, and the k th decision and n th state will be denoted by d_k and s_n , respectively. In the richer “grand world” in which this decision problem is embedded, the DM’s beliefs and preferences are assumed to satisfy the usual axioms of rational choice which guarantee that they are represented by a probability distribution over states and a utility function over consequences.³ Let $p(s)$ denote the DM’s probability for state s , and let $u_k(s)$ denote her utility for the consequence received when decision d_k is made and state s occurs, and define

$$u_{kj}(s) \equiv u_k(s) - u_j(s)$$

as the difference in utilities yielded by decisions d_k and d_j in state s . In these terms, d_k is weakly (strictly) preferred to decision d_j if and only if

$$\sum_{s \in S} p(s) u_{kj}(s) \geq [>] 0. \quad (1)$$

Although we assume that the DM satisfies (or at least wishes to satisfy) the consistency axioms needed to support this representation, we do not assume that she intuitively perceives her complete preference ordering on D or all the parameters of her probability distribution and utility function. Rather, a process of decision analysis is required to help the DM uncover (or perhaps construct) this representation of preferences by focusing myopically on only a few states and/or consequences at a time. Normally this process is carried out by asking the DM to consider various kinds of auxiliary decisions—e.g., choices among gambles or trade-offs among consequences—which are simpler in structure. The DM’s preferences in these auxiliary decisions are used to assess her probabilities and utilities, which are then applied to the solution of the original problem

In standard decision-analytic practice, the auxiliary decisions are merely hypothetical: they have no immediate material consequences for the DM or any other agents. For example, the DM might be asked whether she would be willing to exchange consequence X for a 50-50 roulette-wheel lottery between consequences Y and Z , even though

the opportunity to make this substitution does not exist in the “real” grand world in which the decision problem is embedded.⁴ While this mental exercise may be helpful to the DM in clarifying her own thoughts, it is not typical of the kind of choices that agents ordinarily observe each other to make. For this reason, the probabilities and utilities that a decision analyst would elicit may not be apparent to other agents.

What agents normally do observe are each others’ material choices—i.e., choices which lead to feasible state-contingent transfers of money, commodities, or services between agents. But if we restrict ourselves to material observations, we face the problem cited in the introduction: it may be impossible to uniquely separate probability from utility, because prior stakes are unknown or utilities are otherwise state-dependent. In particular, if the DM’s preferences are represented by probabilities \mathbf{p} and utilities $\{\mathbf{u}_k\}$,⁵ then for any strictly positive function \mathbf{v} , they are equally well represented by probabilities $\mathbf{p}' \propto \mathbf{p}\mathbf{v}$ and utilities $\mathbf{u}'_k \propto \mathbf{u}_k/\mathbf{v}$. The corresponding utility differences satisfy $\mathbf{u}'_{kj} \propto \mathbf{u}_{kj}/\mathbf{v}$, whence $\mathbf{p}\mathbf{u}_{kj} \propto \mathbf{p}'\mathbf{u}'_{kj}$, and consequently:

$$\sum_{s \in S} p(s) u_{kj}(s) \geq [>] 0 \Leftrightarrow \sum_{s \in S} p'(s) u'_{kj}(s) \geq [>] 0.$$

If all that we observe are choices between pairs of decisions, we cannot distinguish the representation $(\mathbf{p}, \{\mathbf{u}_k\})$ from the alternative representation $(\mathbf{p}', \{\mathbf{u}'_k\})$. The same indeterminacy surrounds auxiliary decisions used to estimate probabilities and utilities if the auxiliary decisions involve material rewards.

The simplest kinds of material auxiliary decisions which illustrate this phenomenon are monetary side-gambles pegged to outcomes of the original decision problem. Such gambles are formally similar to transactions in contingent claim markets, where “market makers” set buying and selling prices for contingent claims (e.g., stock options) in order to equalize supply and demand among investors, and investors appear “rational” if the market admits no arbitrage opportunities. In this article, monetary side-gambles will be used as the basic tool for measuring beliefs and preferences, because they resemble the transactions actually observed in the most highly organized markets, and because they reduce the decision problem to its barest essentials: acts of the DM and acts of nature. By focusing entirely on monetary payoffs pegged to these acts, we can gather information about utilities and probabilities without having to catalog the various other consequences of decisions and measure their attributes which are relevant to the DM.⁶ In this framework, the DM will play the role of the (single) investor in the market, and the decision analyst will play the role of the market maker who sets prices such that the investor is indifferent to buying or selling. The purchase or sale of a contingent claim at a specified price is a net gamble, and henceforth we shall refer simply to an “acceptable gamble.” The monetary gambles accepted by the DM constitute the “shadow” that her decision problem casts in a contingent claim market. We will ask what can be inferred about the parameters and the solution of his/her decision problem by observing only this shadow.

To formalize the gambling system, consider a slightly grander-world view of the DM’s decision problem, in which the consequence of decision d_k in state s may be augmented by the receipt of a quantity of money z . A gamble $\mathbf{g} : s \mapsto \mathcal{R}$ is a function that assigns a quantity of money to every state. The set of all gambles is denoted G , and we are interested in the DM’s preferences among decision-gamble pairs.

Assumption 1: The DM's preferences on the set $D \times G$ are represented by a probability distribution $p(s)$ and utility function $u_k(s, z)$, according to which the decision-gamble pair (d_k, \mathbf{g}) is [strictly] preferred to another decision-gamble pair (d_j, \mathbf{g}') if and only if

$$\sum_{s \in S} p(s)(u_k(s, g(s)) - u_j(s, g'(s))) \geq [>] 0.$$

Furthermore, preferences for money are monotonic and smooth, so that $u_k(s, z)$ is a differentiable and strictly increasing function of the monetary payment z for any decision d_k and state s . In the absence of gambling, decision d_k will be chosen only if $(d_k, 0)$ is preferred to $(d_j, 0)$ for every other decision d_j .

Let $u_k \equiv u_k(s, 0)$ and $u_{kj}(s) \equiv u_k(s) - u_j(s)$, so that the utility functions $\{u_k\}$ and the utility differences $\{u_{kj}\}$ continue to represent the DM's preferences among consequences in the absence of gambling, and let $v_k(s)$ denote the DM's *marginal utility for money* when decision d_k is chosen and state s occurs. That is:

$$v_k(s) = \frac{\partial}{\partial z} u_k(s, z) \big|_{z=0}.$$

Then, to a first-order approximation:

$$u_k(s, z) = u_k(s) + zv_k(s) + o(z)$$

as $z \rightarrow 0$, with $v_k(s) > 0$ for all k and s . (Here, $o(z)$ denotes a term with the property that $o(z)/z \rightarrow 0$ as $z \rightarrow 0$.)

A gamble will always be paired with (conditioned on) a decision, and the notation g_k henceforth will be used to denote a gamble paired with decision d_k : the gamble's (first) subscript indicates the conditioning decision. In these terms, the acceptability of gambles is now defined:

Definition: A gamble \mathbf{g}_k is ϵ -acceptable to the DM if, when she chooses d_k , for every $\epsilon > 0$ there exists $\delta > 0$ such that $(d_k, \alpha(\mathbf{g} + \epsilon))$ is strictly preferred to $(d_k, 0)$ for any α between 0 and δ .

In other words, \mathbf{g}_k is ϵ -acceptable if, when "sweetened" by an arbitrarily small amount (ϵ), it is "palatable" to the DM in sufficiently small doses (all $\alpha < \delta$) in the event that she chooses d_k . This ϵ -qualification of acceptability will be used to reconcile the Dutch book argument with nonlinear utility for money.

Lemma 1: \mathbf{g}_k is ϵ -acceptable if it has *nonnegative expected marginal utility* in the event that d_k is chosen, i.e., if

$$\sum_{s \in S} p(s)v_k(s)g_k(s) \geq 0. \tag{2}$$

Proof: In terms of the DM's probabilities and utilities, \mathbf{g}_k is ϵ -acceptable if:

$$\begin{aligned}
0 &< \sum_{s \in S} p(s) [u_k(s, \alpha(g_k(s) + \epsilon)) - u_k(s, 0)] \\
&= \alpha \sum_{s \in S} p(s) v_k(s) (g_k(s) + \epsilon) + o(\alpha) \\
&= \alpha \sum_{s \in S} p(s) v_k(s) g_k(s) + \alpha \epsilon \sum_{s \in S} p(s) v_k(s) + o(\alpha)
\end{aligned}$$

as $\alpha \rightarrow 0$ for any $\epsilon > 0$. The second term on the RHS is strictly positive, while the third term is negligible for sufficiently small α . Hence, the inequality holds in the limit whenever the first term on the RHS is nonnegative. \square

Some other useful properties of ϵ -acceptable gambles are summarized in:

Lemma 2: ϵ -acceptable gambles satisfy:

- (i) Linearity: \mathbf{g}_k is ϵ -acceptable if and only if $\alpha \mathbf{g}_k$ is ϵ -acceptable for any $\alpha > 0$.
- (ii) Additivity: If \mathbf{g}_k and \mathbf{g}'_k are both ϵ -acceptable, then $\mathbf{g}_k + \mathbf{g}'_k$ is ϵ -acceptable.
- (iii) Monotonicity: If \mathbf{g}_k is ϵ -acceptable and $g'_k(s) \geq g_k(s)$ for all s , then \mathbf{g}'_k is ϵ -acceptable.

These properties were taken as primitive by de Finetti (1937, 1974). Here, linearity is a consequence of the ϵ -qualification of gamble acceptance, additivity is a consequence of Lemma 1, and monotonicity is implied by Assumption 1.

3. Probabilities and utilities “revealed” by belief and preference gambles

Thus far it has been assumed that the DM’s preferences in the original decision problem and her acceptance of small monetary side-gambles are jointly represented by a probability distribution \mathbf{p} , utility-differences $\{\mathbf{u}_{kj}\}$, and marginal utilities for money $\{\mathbf{v}_k\}$. Unfortunately, for the reasons pointed out by Kadane and Winkler (1988) and Schervish et al. (1990), neither gamble-acceptance nor any other auxiliary choices among material rewards allow us to uniquely determine \mathbf{p} and $\{\mathbf{u}_{kj}\}$. However, let the DM’s marginal utilities for money be used to define *transformed* probabilities and utility-differences, denoted $\{\hat{\mathbf{p}}_k\}$ and $\{\hat{\mathbf{u}}_{kj}\}$, in the following way:

$$\hat{p}_k(s) \equiv \frac{p(s)v_k(s)}{\sum_{s' \in S} p(s')v_k(s')} \propto p(s)v_k(s), \quad (3)$$

and

$$\hat{u}_{kj}(s) \equiv u_{kj}(s)/v_k(s). \quad (4)$$

Note that, since the marginal utilities are assumed to be strictly positive, $\hat{\mathbf{p}}_k$ is a probability distribution on S for every decision d_k . In these terms,

$$\sum_{s \in S} p(s)u_{kj}(s) > [<, =] 0 \Leftrightarrow \sum_{s \in S} \hat{p}_k(s)\hat{u}_{kj}(s) > [<, =] 0. \quad (5)$$

Thus, the expected value of the DM's true utility difference between decisions d_k and d_j with respect to her true probabilities has the same sign as the expected value of the transformed utility-difference with respect to the transformed probabilities, so that the transformed probabilities and utility-differences also represent the DM's preferences in the original decision problem. Knowing $\{\hat{\mathbf{p}}_k\}$ and $\{\hat{\mathbf{u}}_{kj}\}$ is therefore equivalent to knowing \mathbf{p} and $\{\mathbf{u}_{kj}\}$ for purposes of decision analysis.

The motivation for making this transformation is that $\{\hat{\mathbf{p}}_k\}$ and $\{\hat{\mathbf{u}}_{kj}\}$, unlike \mathbf{p} and $\{\mathbf{u}_{kj}\}$, are uniquely determined⁷ and can be directly revealed through ϵ -acceptable gambles, as the following two lemmas illustrate. First, note that since $\hat{\mathbf{p}}_k$ is a probability distribution on S for any decision d_k , conditional and unconditional probabilities may be defined for all events in the usual way:

$$\hat{p}_k(\mathbf{E}) \equiv \sum_{s \in S} \hat{p}_k(s)E(s), \text{ and}$$

$$\hat{p}_k(\mathbf{E}|\mathbf{F}) \equiv \hat{p}_k(\mathbf{EF})/\hat{p}_k(\mathbf{F}) \text{ if } \hat{p}_k(\mathbf{F}) > 0,$$

for any $\mathbf{E}, \mathbf{F} \subseteq S$. Here, \mathbf{E} and \mathbf{F} are used to denote indicator vectors as well as names for events. Thus, $E(s) = 1$ if event \mathbf{E} includes state s and $E(s) = 0$ otherwise. In these terms, if \hat{p} is a number between 0 and 1, $(\mathbf{E} - \hat{p})\mathbf{F}$ is the net payoff vector of a gamble in which $\$ \hat{p}$ is paid for a lottery ticket that returns \$1 if \mathbf{E} and \mathbf{F} both occur, returns \$0 if \mathbf{F} occurs without \mathbf{E} , and returns the purchase price (i.e., the gamble is called off) if \mathbf{F} fails to occur. This is a conditional gamble on \mathbf{E} given \mathbf{F} at odds of $(1 - \hat{p})/\hat{p}$ against \mathbf{E} . On the probability side of the ledger, we now have:

Lemma 3: For any decision d_k and events \mathbf{E}, \mathbf{F} , the gamble $(\mathbf{E} - \hat{p})\mathbf{F}$ and its negative $(\hat{p} - \mathbf{E})\mathbf{F}$ are both ϵ -acceptable given d_k when $\hat{p} = \hat{p}_k(\mathbf{E}|\mathbf{F})$.

Proof: The expected marginal utility of the first gamble is:

$$\sum_{s \in S} p(s)v_k(s)(E(s) - \hat{p}_k(\mathbf{E}|\mathbf{F}))F(s)$$

$$\propto \hat{p}_k(\mathbf{EF}) - \hat{p}_k(\mathbf{E}|\mathbf{F})\hat{p}_k(\mathbf{F}).$$

The RHS is zero because either $\hat{p}_k(\mathbf{F}) = 0$ (in which case $\hat{p}_k(\mathbf{EF})$ is also zero) or else $\hat{p}_k(\mathbf{EF}) = \hat{p}_k(\mathbf{E}|\mathbf{F})\hat{p}_k(\mathbf{F})$, whence by Lemma 1, both this gamble and its negative are ϵ -acceptable. \square

An ϵ -acceptable gamble of the form $\pm(\mathbf{E} - \hat{p})\mathbf{F}$, conditioned on a decision, will henceforth be referred to as a *belief gamble*, since it reveals information about the DM's beliefs as well as her marginal utilities for money. The quantities $\hat{p}_k(\mathbf{E}|\mathbf{F})$ will be referred to as the DM's "revealed" probabilities for events.

Since by (3) the revealed probabilities are proportional to true probabilities multiplied by marginal utilities for money, Lemma 1 can be rewritten to say that \mathbf{g}_k is ϵ -acceptable if

$$\sum_{s \in S} \hat{p}_k(s) g_k(s) \geq 0. \quad (6)$$

In other words, \mathbf{g}_k is ϵ -acceptable if it has nonnegative *expected monetary value* when expectation is taken with respect to the revealed probabilities $\hat{\mathbf{p}}_k$ rather than the true probabilities \mathbf{p} . The acceptance of such gambles reveals that the DM is seemingly risk neutral but holds (decision-dependent) beliefs $\{\hat{\mathbf{p}}_k\}$. For this reason, the revealed probabilities are also known as “risk neutral probabilities.”⁸

On the utility side of the ledger, the transformed utility-difference vectors are themselves ϵ -acceptable gambles:

Lemma 4: The vector $\hat{\mathbf{u}}_{kj}$ defined by (4) is an ϵ -acceptable gamble conditional on d_k for every k and j .

Proof: d_k will be chosen only if it maximizes subjective expected utility, in which case the following holds for every other decision d_j :

$$0 \leq \sum_{s \in S} p(s) u_{kj}(s) \propto \sum_{s \in S} \hat{p}_k(s) \hat{u}_{kj}(s).$$

Hence, in the event that d_k is chosen, $\hat{\mathbf{u}}_{kj}$ has nonnegative expected monetary value with respect to the DM’s risk neutral distribution, and by Lemma 1, as reformulated in (6), it is ϵ -acceptable. \square

This result can be explained as follows: in the event that d_k is chosen, the gamble $\hat{\mathbf{u}}_{kj}$ yields state-dependent monetary payoffs proportional to $(\mathbf{u}_k - \mathbf{u}_j)/v_k$. These, in turn, provide increments of utility proportional to $\mathbf{u}_k - \mathbf{u}_j$, the difference in utility between decisions d_k and d_j . Therefore, the *expected* increment in utility yielded by $\hat{\mathbf{u}}_{kj}$ when d_k is chosen is nonnegative precisely when the expected utility-difference between decisions d_k and d_j is nonnegative. In this way the DM’s acceptance of $\hat{\mathbf{u}}_{kj}$ merely “affirms” the fact that the expected utility of d_k must be greater than or equal to that of d_j in the event that she chooses d_k . The ϵ -acceptable gambles $\{\hat{\mathbf{u}}_{kj}\}$ will henceforth be called “preference” gambles: they depend *only* on the DM’s preferences for consequences—including money—not on her beliefs. The quantity $\hat{u}_{kj}(s)$ will henceforth be called the DM’s “revealed” utility difference between d_k and d_j in state s , and in these terms the equivalence (5) can be interpreted to say that d_k is preferred to d_j if and only if the expected *revealed* utility-difference d_k between d_j is nonnegative, where the expectation is with respect to the DM’s *revealed* probabilities.

4. Elicitation of belief and preference gambles

The preceding section established that there are two kinds of ϵ -acceptable gambles: “belief” gambles which are determined by the DM’s probabilities (and marginal utilities for money) and “preference” gambles which are determined by utility differences (and marginal utilities for money). Both types of gambles have the same structure: they are

lotteries over states, conditioned on decisions. As such, it may not be apparent how to distinguish one type of gamble from the other in practice. This structural similarity is not necessarily a problem for decision analysis and economic modeling: it will be shown in section 5 that we can infer the expected-utility-maximizing decision merely by observing a sufficiently rich set of ϵ -acceptable gambles, without labeling some as belief gambles and others as preference gambles.

However, in the spirit of divide-and-conquer, it is desirable to have some method of elicitation which enables us not only to distinguish belief gambles from preference gambles, but also to construct them from judgments which are highly decomposed—i.e., which require the contemplation of only a few states and/or decisions at one time. A decomposed method of elicitation is suggested by the following observation: belief gambles are those that would remain ϵ -acceptable if the DM lost her freedom of choice (but meanwhile received no new information), whereas preference gambles are those that would remain ϵ -acceptable if the DM received new information (but meanwhile retained freedom of choice). Hence, we can elicit one type of gamble separately from the other by asking the DM to imagine scenarios in which either her freedom of choice is arbitrarily restricted *or* her information is arbitrarily updated.

Belief gambles are—by definition—those gambles whose acceptability does not depend on the DM's utilities for the original consequences of decisions. Hence, they should remain ϵ -acceptable even if a decision is chosen in some arbitrary and uninformative way. Such gambles can be constructed in a decomposed fashion by asking questions of the following kind: "Suppose that decision d_k is chosen arbitrarily, and consider events **E** and **F** (subsets of the set of states of nature). For what value of \hat{p} would you bet indifferently on or against the occurrence of **E** at odds of $(1 - \hat{p})/\hat{p}$ against **E**,⁹ all conditional on the occurrence of **F**?" The result is $\hat{p}_k(\mathbf{E}|\mathbf{F})$: the DM's revealed probability for **E** | **F** in the event that decision d_k is chosen. By fixing the choice of decision d_k and asking this question for different events **E** and **F**, the DM's revealed probability distribution $\hat{\mathbf{p}}_k$ can be constructed in the manner envisioned by de Finetti.

Preference gambles are—by definition—those gambles which are ϵ -acceptable regardless of the DM's probability distribution over states. Hence, they should remain ϵ -acceptable regardless of any updating of the DM's probabilities prior to the choice of a decision. Such gambles can be elicited in a decomposed fashion by focusing on two decisions and two states at a time—say, decisions d_k and d_j and states s_n and s_m . Assume that d_k is strictly preferred when s_n occurs, while d_j is strictly preferred when s_m occurs—i.e., $u_{kj}(s_n) > 0$ and $u_{kj}(s_m) < 0$ —and suppose it is learned that the true state is one of the two states s_n and s_m , but not which one.¹⁰ Then choosing d_k rather than d_j is tantamount to betting on the true state to be s_n rather than s_m . Let p and $1 - p$ denote the DM's true posterior probabilities¹¹ for states s_n and s_m , respectively. Then d_k will be chosen over d_j precisely when

$$\begin{aligned}
 & pu_k(s_n) + (1 - p)u_k(s_m) \geq pu_j(s_n) + (1 - p)u_j(s_m) \\
 & \Leftrightarrow pu_{kj}(s_n) + (1 - p)u_{kj}(s_m) \geq 0 \\
 & \Leftrightarrow \frac{u_{kj}(s_n)}{-u_{kj}(s_m)} \geq \frac{1 - p}{p}.
 \end{aligned} \tag{7}$$

In other words, d_k will be chosen over d_j if the DM's true posterior odds against s_n , namely, $(1 - p)/p$, are less than or equal to the ratio of the true utility "gain" when s_n occurs to the true utility "loss" when s_m occurs, namely, $u_{kj}(s_n)/(-u_{kj}(s_m))$. It seems as though we could elicit the ratio of true utility-differences by simply asking the DM: "what is the greatest value of the true posterior odds against s_n for which you would choose d_k over d_j ?" Unfortunately, her response would have no material implications: it could not be translated directly into a feasible gamble or trade involving other agents.

However, we *can* ask the DM what the least favorable posterior odds at are which she would ever *bet money on* the occurrence of s_n in the event that she chooses d_k rather than d_j , when only states s_n and s_m are possible. This bet is literally enforceable only if the DM actually learns that the state of nature is in the set $\{s_n, s_m\}$ prior to making his/her decision, but meanwhile it reveals elements of the preference gamble $\hat{\mathbf{u}}_{kj}$, as will now be shown. Let monetary betting odds against s_n be denoted by $(1 - \hat{p})/\hat{p}$. If d_k is chosen, a bet on s_n at those odds is ϵ -acceptable precisely when

$$pv_k(s_n)(1 - \hat{p}) - (1 - p)v_k(s_m)\hat{p} \geq 0$$

$$\Leftrightarrow \frac{1 - \hat{p}}{\hat{p}} \geq \frac{1 - p}{p} \frac{v_k(s_m)}{v_k(s_n)}.$$

In other words, as we might expect, monetary betting odds are ϵ -acceptable if they are greater than or equal to true odds adjusted by the DM's relative marginal utilities for money. Now, from (7), it follows that the largest value which the true odds ratio $(1 - p)/p$ can assume when d_k is chosen is $u_{kj}(s_n)/(-u_{kj}(s_m))$, so that *regardless of the DM's posterior probabilities*, a bet on s_n when d_k is chosen is ϵ -acceptable if the odds satisfy

$$\frac{1 - \hat{p}}{\hat{p}} \geq \frac{u_{kj}(s_n)}{-u_{kj}(s_m)} \frac{v_k(s_m)}{v_k(s_n)} = \frac{\hat{u}_{kj}(s_n)}{-\hat{u}_{kj}(s_m)}.$$

Hence, the minimum odds at which the DM would ever bet on s_n (versus s_m) in the event that she chooses d_k (versus d_j) is simply the ratio of the magnitudes of the "revealed" utility-differences $\hat{u}_{kj}(s_n)$ and $\hat{u}_{kj}(s_m)$ that are elements of the preference gamble $\hat{\mathbf{u}}_{kj}$. By fixing the decisions d_k and d_j and then eliciting these betting odds for different pairs of states s_n and s_m , we can construct the entire gamble $\hat{\mathbf{u}}_{kj}$.

The procedure just described for eliciting preference gambles is admittedly not psychologically trivial. Nevertheless, it is highly decomposed: it requires the DM to contemplate only two decisions and two states at one time, and it depends only on her utilities for the original consequences and marginal utilities for money, not on her probabilities. Furthermore, the response scale (money) is materially significant to the DM, the process does not require the use of artificial randomization devices, and its outcome is subject to immediate empirical validation. The constructed gamble $\hat{\mathbf{u}}_{kj}$ should be holistically perceived as ϵ -acceptable, and belief and preference gambles conditioned on different decisions should satisfy certain consistency conditions implied by Assumption 1. In particular, the pair of preference gambles $\hat{\mathbf{u}}_{kj}$ and $\hat{\mathbf{u}}_{jk}$ and corresponding risk neutral distributions $\hat{\mathbf{p}}_k$ and $\hat{\mathbf{p}}_j$ should be related by:

$$\hat{\mathbf{p}}_k \hat{\mathbf{u}}_j \propto -\hat{\mathbf{p}}_j \hat{\mathbf{u}}_{jk}, \quad (8)$$

and for any three decisions $\{d_i, d_j, d_k\}$, a chain of preference gambles should satisfy:

$$\hat{\mathbf{p}}_k \hat{\mathbf{u}}_{kj} \propto \beta \hat{\mathbf{p}}_k \hat{\mathbf{u}}_{ki} + (1 - \beta) \hat{\mathbf{p}}_i \hat{\mathbf{u}}_{ij} \quad (9)$$

for some $\beta \in (0, 1)$. These relations all follow from the fact that $\hat{\mathbf{p}}_k \hat{\mathbf{u}}_{kj} \propto \mathbf{p}(\mathbf{u}_k - \mathbf{u}_j)$.

5. The coherence theorem

Once the DM's revealed probabilities $\{\hat{\mathbf{p}}_k\}$ and revealed utility-differences $\{\hat{\mathbf{u}}_{kj}\}$ have been elicited by the procedures described in the preceding section, it follows from (5) that we can determine the optimal decision by evaluating $\sum_{s \in S} \hat{p}(s) \hat{u}_{kj}(s)$ for every pair of decisions d_k and d_j , concluding that d_k is weakly (strictly) preferred to d_j if this quantity is nonnegative (positive). Thus, by observing the DM's choices in the small world of monetary gambles, information is obtained which suffices to prescribe her choice in the grand(er) world of the original decision problem. But more can be said about the connection between grand-world and small-world behavior, namely, that rational behavior in the grand world, as codified in the axioms of subjective expected-utility theory, corresponds to rational behavior in the small world, as codified in principles of "coherent" gambling.

Recall that the belief gambles conditioned on decision d_k reveal what the DM's risk neutral probabilities actually are when this decision is chosen, namely $\hat{\mathbf{p}}_k$, whereas the preference gambles conditioned on d_k reveal constraints that the risk neutral probabilities must satisfy if this decision is to be chosen rationally, namely, $\sum_{s \in S} \hat{p}(s) \hat{u}_{kj}(s) \geq 0$ for every other decision d_j . It will now be shown that if these constraints are not satisfied by the actual risk neutral probabilities, then the belief and preference gambles conditioned on d_k lead to a Dutch book—i.e., arbitrage or sure loss—a condition which de Finetti termed *incoherence*. Thus, a DM who failed to maximize her revealed subjective expected utility would be an incoherent gambler.

The concept of coherence originally defined by de Finetti is weak, or *ex ante*, coherence, which forbids the DM to accept a sure loss—i.e., a strictly negative gamble. A stronger requirement is *strict* coherence, which forbids the DM to accept any possibility of loss without some compensating possibility of gain—i.e., it forbids her to accept even a *semi-negative* gamble. Strict coherence requires the DM to implicitly assign positive probability to every logically possible event, ruling out weakly dominated acts. The standard which will be applied here is intermediate in strength between weak and strict coherence:

Definition: The decision maker is *ex post coherent* in outcome (d_k, s_n) if there does not exist an ϵ -acceptable gamble \mathbf{g}_k conditioned on d_k such that $g_k(s) \leq 0$ for all $s \in S$ and $g_k(s_n) < 0$.

As the name implies, *ex post* coherence is resolved after the outcome of the problem is known: it will turn out to require the DM to have implicitly assigned positive probability to the outcome which was observed, but not necessarily to those that were not. The

rationale for this compromise between ex ante and strict coherence is that the set of outcomes here involves decisions as well as states of nature, and it is desirable to treat them symmetrically from the perspective of a betting opponent who may not know or care which events are under the DM's control. If the opponent succeeds in extracting money from the DM without having put a proportional amount of money at risk, then the DM appears irrational after-the-fact, regardless of whether the loss was precipitated by an unexpected event or an unwise decision: she should have known better *or else* acted otherwise.

Theorem: The DM is ex post coherent in outcome (d_k, s_n) if and only if d_k maximizes expected utility and s_n is assigned positive probability.

Proof: Let G_k denote the matrix whose columns are the ϵ -acceptable belief and preference gambles conditioned on d_k . By the properties of linearity, additivity, and monotonicity summarized in Lemma 2, any gamble which equals or dominates $G_k \alpha$ is also ϵ -acceptable, for any nonnegative vector α . Then the DM is ex post coherent in outcome (d_k, s_n) if and only if there is no $\alpha \geq 0$ such that $G_k \alpha \leq 0$ with strict inequality in state s_n . By linear duality (e.g., Gale, 1960, Theorem 2.10), this is true if and only if there exists a distribution π such that $\pi^T G_k \geq 0$ and $\pi(s_n) > 0$, i.e., a distribution assigning nonnegative expectation to every ϵ -acceptable gamble conditioned on d_k and assigning positive probability to s_n . The belief gambles conditioned on d_k have nonnegative expectation under π if and only if $\pi = \hat{p}_k \propto p v_k$ for all s ; and $\pi(s_n) > 0$ implies $p(s_n) > 0$ since $v_k(s) > 0$ for all s . The preference gambles conditioned on d_k have nonnegative expected value under π if and only if, for every other decision d_j ,

$$\begin{aligned} \sum_{s \in S} \pi(s) \hat{u}_{jk}(s) &\geq 0 \\ \Leftrightarrow \sum_{s \in S} \hat{p}_k(s) \hat{u}_{jk}(s) &\geq 0 \\ \Leftrightarrow \sum_{s \in S} p_k(s) u_{jk}(s) &\geq 0. \end{aligned}$$

That is, d_k has greater expected utility than d_j . □

This result suggests a novel computational approach for determining optimal decisions: having elicited belief and preference gambles conditional on d_k , invoke the properties of linearity, additivity, and monotonicity, then use linear programming to determine whether there is any state s_n such that the DM is ex post *coherent* in outcome (d_k, s_n) . In other words, form the matrix G_k whose columns are the payoff vectors of ϵ -acceptable belief and preference gambles conditioned on d_k , and solve for a vector $\pi \geq 0$, satisfying $\pi^T G_k \geq 0$ and $\sum_{s \in S} \pi(s) = 1$, which is a probability distribution assigning nonnegative expectation to every ϵ -acceptable gamble. If such a distribution exists, then the DM is ex post coherent in those outcomes $\{(d_k, s_n)\}$ for which $\pi(s_n) > 0$, and consequently d_k is a potentially optimal decision. If such a distribution does not exist, then, by duality, there exists instead a nonnegative vector α such that $G_k \alpha < 0$ —i.e., a weighted combination of the gambles yielding a sure loss conditional on decision d_k —and the decision evidently is suboptimal.

Note that the DM has not been asked to assign probabilities to his/her own decisions: decisions have been treated as “events” only for purposes of conditioning. However, the acceptance of a sure loss conditional on a decision is operationally equivalent to assigning zero probability to that decision. Technically this is a risk neutral probability, but there is no distinction between true and risk neutral probabilities when both are zero. So the theorem has the following implication: *by eliciting a sufficiently rich set of ϵ -acceptable gambles and invoking the properties of linearity, additivity, and monotonicity, it is discovered that the DM implicitly assigns zero probability to decisions which ultimately fail to maximize subjective expected utility.*

Thus, coherence (together with linearity, etc.) is what is observed when expected-utility-maximizing grand-world behavior is projected into the small world of money-based measurements. This result establishes that conformity with subjective expected-utility theory is a sufficient as well as a necessary condition for the avoidance of Dutch books *even in the presence of nonlinear utility for money*, a point which has sometimes been disputed (e.g., Yaari, 1985). The avoidance of Dutch books ex post as well as ex ante also requires the implicit assignment of positive probability to the state which was observed, if not to every state.

6. An example of decision analysis

The elicitation and optimization procedures developed in the preceding sections are intended to illustrate the scope of “material” decision analysis, rather than to supplant the psychological techniques commonly used in practice (c.f. von Winterfeldt and Edwards, 1986). Nevertheless, it is instructive to consider how, in principle, a divide-and-conquer decision analysis might be carried out in terms of the DM’s revealed (risk neutral) probabilities $\{\hat{p}_k\}$ and revealed utility-differences $\{\hat{u}_{kj}\}$. The following simplified example illustrates this process.

Several months before harvest time, a wheat farmer has the opportunity to sell part or all of her crop “forward” at the current price of \$3.80 per bushel. That is, the farmer can contract now to deliver a specified quantity of wheat in the future for a guaranteed price of \$3.80. Suppose that he/she is considering one of three options: selling forward 100%, 60%, or 0% of the expected amount of his/her crop. Any portion of the crop not sold forward will be sold at the spot price prevailing at the time of harvest, which is expected to take on one of three values: \$3, \$4, or \$5 per bushel. (These three-fold discretizations of the decision space and state space are small-world simplifications of the farmer’s problem, and are typical of much applied decision analysis. A more detailed analysis of a similar problem is given by Hildreth, 1979). The decision tree for the farmer’s problem is shown in figure 1. The values attached to the terminal branches are the average prices obtained per bushel, which are shown merely to indicate an approximate ranking of outcomes. For example, if the farmer sells only 60% of the crop forward, and the spot price turns out to be \$3 per bushel, then the average price per bushel is $0.6 \times \$3.80 + 0.4 \times \$3.00 = \$3.48$.

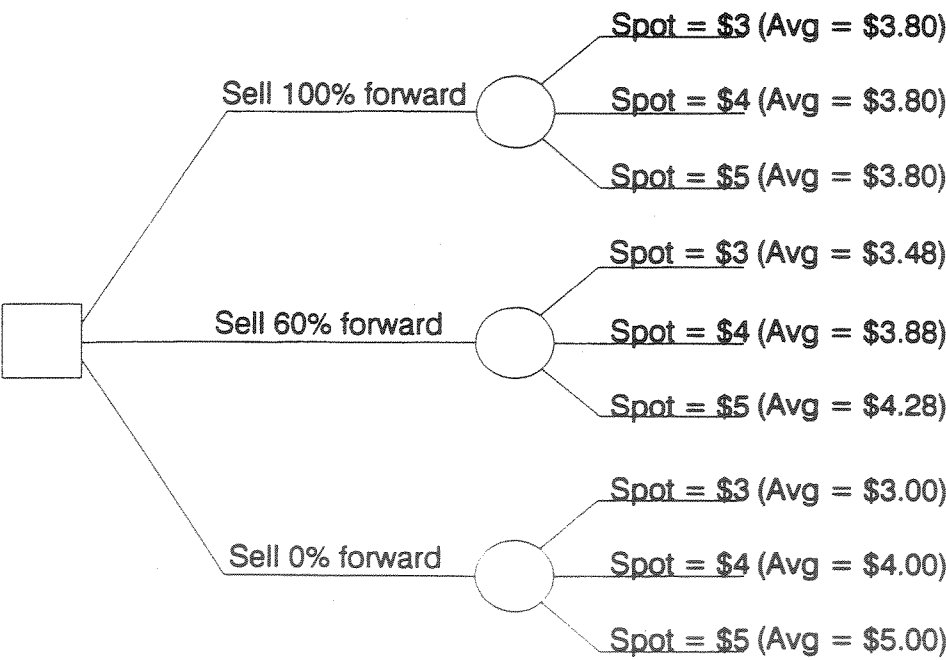


Figure 1. Wheat farmer's decision tree.

To carry out the decision analysis, suppose the farmer is first asked to imagine that decision d_1 has been chosen (100% sold forward), and to determine the prices at which she would indifferently buy or sell lottery tickets paying \$1 in the event that states s_1, s_2 , or s_3 occur (spot prices of \$3, \$4, or \$5, respectively). Suppose that she assesses these prices as \$0.30, \$0.50, and \$0.20. (The prices must sum to 1 or else arbitrage is possible.) Then her risk neutral probability distribution given d_1 is $\hat{p}_1 = (0.3, 0.5, 0.2)$. Since d_1 is the decision yielding the most nearly “constant” outcome, \hat{p}_1 might be considered to approximate her true distribution, although this is not necessary for the analysis. Now suppose that the same exercise is carried out under the assumption that d_2 has been chosen (60% sold forward), yielding $\hat{p}_2 = (0.35, 0.5, 0.15)$. The risk neutral probability of s_1 is greater when d_2 is chosen than when d_1 is chosen (0.35 versus 0.3), apparently reflecting a higher marginal utility for money under a lower state of final wealth (\$3.48 received per bushel versus \$3.80). Finally, suppose that on the assumption that d_3 has been chosen (0% sold forward), the farmer assesses $\hat{p}_3 = (0.45, 0.45, 0.10)$.

Now consider the assessment of preference gambles, beginning with \hat{u}_{12} . Note that d_1 (100% sold forward) should be preferred to d_2 (60% sold forward) in state s_1 (\$3 spot price), and vice versa in state s_2 (\$4 spot price). The average prices are \$3.80 versus \$3.48, given the \$3 spot price, and \$3.80 versus \$3.88, given the \$4 spot price. Now let the farmer imagine that she has received information leading him/her to believe that one of these two spot prices will prevail. If she chooses d_1 in this situation, at what odds would she also accept a small bet on the \$3 spot price occurring? Note that this choice yields a “gain” of \$0.32 per bushel if the \$3 spot price occurs, and a “loss” of \$0.08 per bushel if the \$4 spot price occurs. The ratio of monetary gain to loss is 4:1 in the primary decision problem,

but suppose that the farmer judges that she would bet on the \$3 spot price at minimum odds-against of 9:2, reflecting a diminished relative utility for wealth in state s_1 when she has chosen d_1 . This establishes that $\hat{u}_{12}(s_1)/(-\hat{u}_{12}(s_2)) = 9/2$.

Repeating this process with s_3 replacing s_2 , note that the “loss” for choosing d_1 over d_2 is \$0.48 per bushel when s_3 occurs, as compared to the “gain” of \$0.32 per bushel when s_1 occurs. The ratio of gain to loss is now 2:3, but suppose that the farmer is only willing to be on s_1 at odds-against of 9:11 or better, yielding $\hat{u}_{12}(s_1)/(-\hat{u}_{12}(s_3)) = 9/11$. Altogether, this gives $\hat{u}_{12} \propto (1, -2/9, -11/9) \approx (1.0, -0.22, -1.22)$. Continuing with a few more between-decision comparisons suppose we obtain $\hat{u}_{13} = (1.0, -0.2, -0.9)$ and $\hat{u}_{23} = (1.0, -0.2, -1.1)$. This information suffices to determine:

$$\sum_{s \in S} \hat{p}_1(s) \hat{u}_{12}(s) = (0.3)(1.0) + (0.5)(-0.22) + (0.2)(-1.22) = -0.054 < 0 \Rightarrow d_1 < d_2,$$
$$\sum_{s \in S} \hat{p}_1(s) \hat{u}_{13}(s) = (0.3)(1.0) + (0.5)(-0.2) + (0.2)(-0.9) = 0.02 > 0 \Rightarrow d_1 > d_3,$$
$$\sum_{s \in S} \hat{p}_2(s) \hat{u}_{23}(s) = (0.35)(1.0) + (0.5)(-0.2) + (0.15)(-1.1) = 0.085 > 0 \Rightarrow d_2 > d_3,$$

from which it may be concluded that $d_2 > d_1 > d_3$ —i.e., sell 60% > sell 100% > sell 0%. Notice that \hat{p}_3 , \hat{u}_{21} , \hat{u}_{31} , and \hat{u}_{32} are not needed to determine this preference ordering—assuming that everything else has been assessed accurately—but these would be useful for cross-checking and reconciling inconsistencies via equations (8) and (9). \hat{p}_1 , \hat{p}_2 , \hat{u}_{12} , \hat{u}_{23} , and \hat{u}_{13} already satisfy (9) within reasonable limits with $\beta = 0.39$.¹²

To illustrate the alternative method developed in section 5 for determining the optimal decision, form the matrix G_1 , whose columns are the belief and preference gambles conditioned on decision d_1 , as follows:

Table 1. Belief and preference gambles conditioned on d_1

Outcome	$E_1 - \hat{p}_1(E_1)$	$\hat{p}_1(E_1) - E_1$	$E_2 - \hat{p}_1(E_2)$	$\hat{p}_1(E_2) - E_2$	$E_3 - \hat{p}_1(E_3)$	$\hat{p}_1(E_3) - E_3$	\hat{u}_{12}	\hat{u}_{13}
s_1	0.8	-0.8	-0.5	0.5	-0.3	0.3	1.0	1.0
s_2	-0.2	0.2	0.5	-0.5	-0.3	0.3	-0.22	-0.2
s_3	-0.2	0.2	-0.5	0.5	0.7	-0.7	-1.22	-0.09

A sure loss conditional on d_1 is achieved by setting $\alpha = (0, 25, 0, 0, 20, 0, 20, 0)$, yielding $G_1\alpha = (-6, -6, -6)$, proving that d_1 is suboptimal. The same procedure can be repeated for decisions d_2 and d_3 to reveal that d_3 is also suboptimal, whereas d_2 is optimal.

The decision in this problem concerns the amount of a commodity to be bought or sold at a given market price. In such a problem, it is actually unnecessary to elicit preference gambles: it suffices to elicit the DM’s risk neutral probabilities and to compare his/her own “risk neutral valuation” of a unit of the commodity (here, a bushel of wheat) with its market price (here, \$3.80) under each of the different decision scenarios:

Risk neutral valuation given $d_1 = (0.3)(\$3) = (0.5)(\$4) + (0.2)(\$5) = \3.90

Risk neutral valuation given $d_2 = (0.35)(\$3) + (0.5)(\$4) + (0.15)(\$5) = \3.80

Risk neutral valuation given $d_3 = (0.45)(\$3) + (0.45)(\$4) + (0.1)(\$5) = \3.65

It is seen that under d_2 her risk neutral valuation equals the market price, indicating that she would prefer to sell forward neither more nor less wheat than she is committed to sell under this option (60%). Under d_1 her risk neutral valuation is above market price, indicating that she would prefer to sell less than she is committed to under this option (100%), and conversely for d_3 . These relationships illustrate a general principle of decision analysis with risk neutral probabilities: in the vicinity of her optimal decision, the DM's risk neutral valuation for any marketed commodity must equal its market price. A DM with convex (risk averse) preferences can therefore solve a buying/selling problem by searching among the available options for one in which her risk neutral valuation satisfies this first-order condition for an optimum, at which point she is in equilibrium with the market. In the example, if the DM's risk neutral valuation had *not* equalled the market price for any of the three options originally considered, then some percentage other than 0%, 60%, or 100% would have been optimal. Further discussion of this principle is given in Nau and McCardle (1991) and Smith and Nau (1994).

7. Discussion

The inability to measure true personal probabilities via material methods may be less dire for statistical theory than has been suggested by Kadane and Winkler (1988) and Schervish et al. (1990): the fact that a theory relies on a construct which is not directly measurable does not necessarily invalidate the theory as long as it correctly accounts for phenomena which *are* measurable (Feynman et al., 1965, pp. 2–8, 2–9). Besides, in a nondecision (pure inference) context, the distorting effect of marginal utilities for money is confined to the subject's prior probabilities, which she is not required to justify. Bayesian inference emphasizes the process of prior-to-posterior updating, in which the prior is multiplied by a likelihood function summarizing new information. This process is unaffected if the subject's prior distribution is taken to be a risk neutral distribution rather than a true distribution. If the new information per se does not affect the subject's utilities, then her risk neutral probabilities are updated in the same manner (i.e., multiplied by the same likelihood function) as her true probabilities. The multiplicative contribution of her state-dependent marginal utilities for money propagates straight through from prior to posterior. Thus, demanding the separation of probability from utility in an inferential setting is tantamount to demanding that the subject explain the origin of her prior distribution.

In a decision context, the need to know true rather than risk neutral probabilities depends on the information about utilities which is available. If the decision maker's true utility-differences are known, then of course they should be weighted by true probabilities, but this begs the uniqueness question with respect to the measurement of utilities.

The results of this article show that it is always possible, at least in principle, to uniquely measure transformed utility-differences which can be weighted by risk neutral probabilities to determine the decision maker's utility-maximizing decision.

It might be objected that the nominal decision maker is sometimes an "expert" whose personal probabilities are solicited for use by others in their own decision making. In such cases, it appears as though the latter would want the expert's true probabilities, undistorted by her marginal utilities for money. But this begs the question: how does one agent *credibly* and *verifiably* communicate her beliefs to others? How do the recipients of the expert's probability forecast know that they are getting her honest opinion, and to what extent can it be confirmed by an outside observer that the recipients have assimilated the expert's message? If the climate is competitive rather than cooperative, it may be hard to certify the transfer of information unless the participants back up their pronouncements by betting with the others or submitting to some other monetary reward scheme—but monetary transactions reveal only risk neutral probabilities! Furthermore, if such a scheme is actually implemented, the agents will have incentives to continue betting with each other, transferring state-contingent wealth, until a competitive equilibrium is reached in which their risk neutral probabilities are equal. (This scenario is discussed by Kadane and Winkler, 1988, and in more detail by Nau, 1994.) Thus, to all outward appearances, the probabilities which can be communicated between agents are their risk neutral probabilities rather than their true probabilities. If they appear to converge on a common distribution, it must be a risk neutral distribution: their true probabilities may still differ.

By the same token, the utility-differences which can be visibly communicated between agents are not the true utility-differences $\{u_{kj}\}$ but rather the transformed values $\{\hat{u}_{kj}\}$ revealed by preference gambles. This observation is relevant to the modeling of games of strategy, in which the players' utility-differences constitute the "rules of the game" which are supposed to be common knowledge. The implications of using preference gambles to elicit the rules of noncooperative games have been explored by Nau and McCardle (1990) and Nau (1992ab). The small-world manifestation of mutually expected Bayesian rationality in such games turns out to be precisely the property of ex post coherence applied to the *joint* behavior of all the players—i.e., their joint decision-making behavior should present no ex post arbitrage opportunities to an outside betting opponent. The corresponding duality theorem states that the only outcomes in which the players are jointly coherent are those which occur with positive probability in a *correlated equilibrium* of the game, which is a generalized Nash equilibrium in which mixed strategies may be correlated between players (Aumann, 1974, 1987). This result supports Aumann's contention that correlated equilibrium is "the" expression of Bayesian rationality in noncooperative games. It also responds to some of the subjectivist complaints against game theory raised by Kadane and Larkey (1982), albeit under a novel interpretation of common knowledge: the correlated equilibrium distribution, like the common prior, should be interpreted as a risk neutral distribution.

The preceding arguments take a radically operational view of communication: certainly not all public communication of beliefs and preferences is mediated by transfers of wealth. Yet, if the axioms of subjective expected-utility theory are taken seriously as

standards of economic rationality, then the mutual revelation of probabilities and utilities *by whatever means* must either equalize the risk neutral probabilities of the agents involved, or else it must leave them in a state of common knowledge of Pareto inefficiency, which could be alleviated by trading contingent claims. The appealing feature of de Finetti's elicitation method is that, as a microcosm of a contingent claim market, it addresses the agents' desires for efficiency as well as for information.

Acknowledgments

The author is grateful to Greg Fischer, Harvey Gold, Nancy Keeshan, Kevin McCardle, Jim Smith, Peter Wakker, and Bob Winkler for comments on earlier drafts of this article. This research was supported by the Business Associates Fund and the Hanes Corporation Foundation Fund at the Fuqua School of Business.

Notes

1. The focus in this article is on what can be learned about beliefs and preferences through "material" measurements—i.e., measurements which are immediately backed up by gambles or trades involving exchanges of money or commodities with other agents—as opposed to psychological measurements, in which purely intuitive judgments are elicited (e.g., DeGroot, 1970) and/or choices are made among purely hypothetical alternatives.
2. In Savage's (1954) subjective expected utility theory, it is assumed that consequences can always be defined in such a way as to justify the assumption that their "true" utilities are state-independent. However, these definitions of consequences may be elaborate and inconvenient. Schervish et al. (1990) note that outside observers generally cannot determine which consequences have constant utility across states for a particular agent, and that the agent's "indirect" utilities for consequences which are easily measurable are often state-dependent due to unknown prior stakes. It is sometimes convenient to assume that even true utilities are intrinsically state-dependent (Karni et al., 1983; Karni, 1985, 1992; Drèze, 1987), which would not affect the results in this article.
3. It is not necessary for our purposes to specify the consequences or measure their attributes: it suffices to let utility be defined simply as a function of the decision and the state. Hence, rather than beginning with an infinite set of grand-world consequences and invoking Savage's (1954) axioms, we could use a conditional preference model such as Fishburn (1973, 1982, chapter 12) or Krantz et al. (1971, chapter 8), with the stipulation that the states of nature are the same for all decisions.
4. Other examples of purely hypothetical auxiliary decisions are choices between decision-event pairs (Fishburn, 1973; 1982, chapter 12) and choices among lotteries predicated on the assignment of an objective probability distribution to states of nature (Karni et al., 1973; Karni, 1985; Schervish et al., 1990).
5. Boldface notation is used to denote functions of the state s , which can be considered as vectors whose elements are indexed by s . Multiplication of vectors is defined pointwise—e.g., \mathbf{pv} is the vector whose s th element is $p(s)v(s)$.
6. Gambling payoffs are allowed to depend on the DM's own acts as conditioning events in order to articulate the "utility" side of the decision problem. In this respect, the gambles considered here differ from most transactions in real contingent claim markets, which are pegged to events relevant to everyone and controlled by no one—e.g., states of nature corresponding to stock and commodity prices.
7. Uniqueness of transformed probabilities and utilities is relative to the monetary currency: if the unit of currency is changed, then $\{\mathbf{v}_k\}$ and $\{\hat{\mathbf{p}}_k\}$ may change, as in the dollar-yen example of Schervish et al. (1990).

However, $\{\hat{u}_{kj}\}$ changes reciprocally, and the effects cancel out when the products $\hat{p}_k(s)\hat{u}_{kj}(s)$ are taken in (5): the representation of preferences is unaffected.

8. The term "risk neutral probability" originated in the finance literature (Cox and Ross, 1976) where it refers to the apparent beliefs of a risk neutral representative investor in an arbitrage-free securities market. The fact that marginal utilities are impounded in risk neutral probabilities was pointed out by Rubinstein (1976), Garman (1979), and Brennan (1979), and even earlier by Drèze (1970), who observed: "It may come as a surprise to the reader that, while the prices for contingent claims to the numeraire ... have all the properties of a probability measure, they are to be interpreted as the products of a probability by a relative marginal utility."
9. A bet on [against] E at odds of $(1 - \hat{p})/\hat{p}$ against E is a bet which yields a net gain [loss] proportional to $1 - \hat{p}$ if E occurs and a net loss [gain] proportional to \hat{p} if E does not occur. Thus, the higher the odds against E , the more [less] favorable is a bet on [against] E .
10. If there are not two states which yield different strict preferences between decisions d_k and d_j , then evidently one decision at least weakly dominates the other, and it is unnecessary to precisely quantify the utility differences between them.
11. The posterior probabilities p and $1 - p$ could be merely a renormalization of the prior probabilities $p(s_n)$ and $p(s_m)$, or they could be revised in some other way by the receipt of information. It is not necessary to attach specific values to them for the purposes of assessing preference gambles.
12. Of course, in practice the not-so-primitive judgments in the wheat farmer's problem could be aided by more conventional techniques of probability and utility assessment. In fact, the values given here are consistent (up to rounding off) with the assumption of a crop size of 100,000 bushels, a true probability distribution of (0.2, 0.5, 0.3), and a state-independent exponential utility function with a risk tolerance of \$200,000. That is, $u(z) \propto 1 - \exp(-z/200000)$ where z is the total received from wheat sales and gambling. (The risk tolerance is the reciprocal of the Arrow-Pratt measure of absolute risk aversion, $-u''(z)/u'(z)$, which is constant for the exponential utility function.)

References

- Aumann, R. (1974). "Subjectivity and Correlation in Randomized Games," *Journal of Mathematical Economics* 1, 67-96.
- Aumann, R. (1987). "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica* 55, 1-18.
- Brennan, M. (1979). "The Pricing of Contingent Claims in Discrete Time Markets," *Journal of Finance* 24(1), 53-68.
- Cox, J., and S. Ross. (1976). "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics* 3(1/2), 145-166.
- de Finetti, B. (1937). "La Prévision: ses lois logiques, ses sources subjectives," *Annals Institut Henri Poincaré* 7, 1-68. Translation reprinted in 1980 in H.E. Kyburg and H.E. Smokler (eds.), *Studies in Subjective Probability*, 2nd ed. New York: Robert Krieger, 53-118.
- de Finetti, B. (1974). *Theory of Probability*, Vol. 1. New York: Wiley.
- DeGroot, M. (1970). *Optimal Statistical Decisions*. New York: Wiley.
- Drèze, J. (1970). "Market Allocation Under Uncertainty," *European Economic Review* 2, 133-165. Reprinted in Drèze, 1987.
- Drèze, J. (1987). *Essays on Economic Decisions Under Uncertainty*. Cambridge University Press.
- Feynman, R., R. Leighton, and B. Sands. (1965). *Lectures on Physics*, Vol. 3. Reading, MA: Addison-Wesley.
- Fishburn, P.C. (1973). "A Mixture-Set Axiomatization of Conditional Subjective Expected Utility," *Econometrica* 41, 1-25.
- Fishburn, P.C. (1982). *The Foundations of Expected Utility*. Dordrecht: D. Reidel.
- Gale, G. (1960). *The Theory of Linear Economic Models*. New York: McGraw-Hill.
- Garman, M. (1979). "A Synthesis of the Pure Theory of Arbitrage." Working Paper No. 98, Institute of Business and Economic Research, University of California, Berkeley.
- Harsanyi, J. (1967). "Games with Incomplete Information Played by 'Bayesian' Players: I-III," *Management Science* 14, 159-182, 320-334, 486-502.

- Hildreth, C. (1979). "An Expected Utility Model of Grain Storage and Hedging by Farmers." Technical Bulletin 321, Agricultural Experiment Station, University of Minnesota.
- Kadane, J.B., and P.D. Larkey (1982). "Subjective Probability and the Theory of Games," *Management Science* 28(2), 113–125.
- Kadane, J.B., and R.L. Winkler. (1988). "Separating Probability Elicitation from Utilities," *Journal of the American Statistical Association* 83, 357–363.
- Karni, E., D. Schmeidler, and K. Vind. (1983). "On State-Dependent Preferences and Subjective Probabilities," *Econometrica* 51, 1021–1031.
- Karni, E. (1985). *Decision Making Under Uncertainty: The Case of State-Dependent Preferences*. Cambridge: Harvard University Press.
- Karni, E. (1992). "Subjective Probabilities and Utility with Event-Dependent Preferences," *Journal of Risk and Uncertainty* 5(2), 107–125.
- Krantz, D., R. Luce, P. Suppes, and A. Tversky. (1971). *Foundations of Measurement, Volume 1: Additive and Polynomial Representations*. New York: Academic Press.
- Nau, R., and K. McCardle. (1990). "Coherent Behavior in Noncooperative Games," *Journal of Economic Theory* 50(2), 424–444.
- Nau, R., and K. McCardle. (1991). "Arbitrage, Rationality, and Equilibrium," *Theory and Decision* 31, 199–240.
- Nau, R. (1992a). "Joint Coherence in Games of Incomplete Information," *Management Science* 38(3), 374–387.
- Nau, R.F. (1992b). "The Relativity of Subjective and Objective Correlated Equilibria." Working paper, Fuqua School of Business, December 1992.
- Nau, R. (1994). "The Incoherence of Agreeing to Disagree." *Theory and Decision*, forthcoming.
- Ross, S. (1976). "Return, Risk, and Arbitrage," *Risk and Return in Finance*, I. Friend and J. Bicksler, Eds., Cambridge: Ballinger.
- Rubinstein, M. (1976). "The Valuation of Uncertain Income Streams and the Pricing of Options," *Bell Journal of Economics and Management Science* 7, 407–425.
- Savage, L.J. (1954). *The Foundations of Statistics*. New York: Wiley.
- Schervish, M.J., T. Seidenfeld, and J.B. Kadane. (1990). "State-Dependent Utilities," *Journal of the American Statistical Association* 85(411), 840–847.
- Smith, J.E., and R.F. Nau. (1994). "Valuing Risky Projects: Option Pricing Theory and Decision Analysis," *Management Science*, forthcoming.
- Varian, H. (1987). "The Arbitrage Principle in Financial Economics," *Journal of Economic Perspectives* 1, 55–72.
- von Winterfeldt, D., and W. Edwards. (1986). *Decision Analysis and Behavioral Research*. Cambridge University Press.
- Yaari, M. (1985). "On the Role of 'Dutch Books' in the Theory of Choice Under Risk." Nancy Schwarz Memorial Lecture, Northwestern University.