

Subjective Probabilities on a State Space[†]

By EDI KARNI*

This paper extends the analytical framework of Karni (2011) to include a state space and advances a choice-based definition of subjective probabilities on this space. These probabilities represent the beliefs of Bayesian decision makers regarding the likelihoods of events, thus resolving a long-standing, fundamental difficulty with the definition of subjective probabilities. (JEL D81, D83)

Beginning with Frank P. Ramsey (1931) and Bruno de Finetti (1937) and culminating with Leonard J. Savage (1954), the concept of choice-based subjective probabilities has been a subject of much interest and some controversy. With rare exceptions, subjective probabilities are treated as an aspect of the representation of a decision maker's preference relation, which *defines* his degree of belief regarding the likelihood of events. According to this approach, whether, in his mind, the decision maker entertains beliefs that can be quantified by probabilities, is immaterial as long as his preferences may be represented using such probabilities.¹

A radically different approach considers the decision maker's beliefs to be a cognitive phenomenon that feeds into the decision-making process. According to this approach the subjective probabilities *measure*, rather than define, the decision maker's beliefs.

To grasp the difference between the two notions of subjective probability, it is convenient to think of the decision maker as a black box capable, when presented with pairs of alternatives, to select the preferred alternative, or express indifference between the two. Invoking this metaphor, consider the following thought experiment. Let us “construct” a black box into which we upload probabilities of events, utilities of consequences, and a set of instructions to calculate for each alternative its expected utility, and to select the alternative that yields the highest expected utility. Next we present the box with pairs of alternatives and record its choices. Is it possible to recover the probabilities and utilities that were uploaded into the black box by observing its choice pattern?

To answer this question, suppose that a black box is instructed to choose among bets on the outcome of the final game of the world cup, mimicking the betting

* Department of Economics, Johns Hopkins University, 34th and Charles Street, Baltimore, MD 21218 (e-mail address: karni@jhu.edu). I am grateful to Jacques Drèze, Brian Hill, and an anonymous referee for their comments on an earlier draft of this paper and to Efe Ok and Marco Scarcini for enlightening conversations and useful suggestions.

[†] To comment on this article in the online discussion forum, or to view additional materials, visit the article page at <http://www.aeaweb.org/articles.php?doi=10.1257/mic.3.4.172>.

¹ This is the approach taken by Savage (1954) and all subsequent contributions invoking Savage's analytical framework.

behavior of a soccer enthusiast whose country made it to the final. Suppose that our soccer fan believes that his country has equal chance of winning or losing the final match. Accordingly, let the black box assign equal probabilities to victory and defeat. The consequences of the bets are their state-contingent monetary payoffs. Assume that the soccer fan's utility of money depends on the whether his country wins (victory state) or loses (defeat state). To keep things simple, suppose that, for the same payoff, his utility in the victory state is twice that of the defeat state. Consequently, instructed to calculate expected utility, the black box assigns the bet that pays x_v and x_d in the victory and defeat states, respectively, the expected utility $0.5 \times 2u(W + x_v) + 0.5 \times u(W + x_d)$, where u is the utility function uploaded into the black box and W denotes the soccer fan's initial wealth. Clearly, the "true beliefs" of the black box are represented by the uniform probability distribution.

To infer the probabilities from the black box's choices, apply any of the standard elicitation methods. For instance, applying the quadratic scoring rule method, we ask the black box to select the value of α , thereby choose a bet whose payoff is $x_v(\alpha) = -r\alpha^2$ and $x_d(\alpha) = -r(1 - \alpha)^2$, where $r > 0$, in the win and loss state, respectively. It is well known that, as r tends to zero, the optimal value of α is an estimate of the probability of the loss state. In this example, as r tends to zero, the optimal α tends to $1/3$. Clearly, this probability is not what was uploaded into the black box, and consequently, does not represent its true beliefs. In fact, the application of the quadratic scoring rule identified the value of the normalized coefficient of $u(W + x_d)$. In other words, using the output of the scoring rule method, we attribute to the black box the calculation of the expected utility of the bets according to formula $2/3 \times u(W + x_v) + 1/3 \times u(W + x_d)$. Selection of the bets according to this formula is exactly the same as the choice behavior governed by the original set of instructions. The probabilities $2/3$ and $1/3$ *define*, rather than measure, the black box's beliefs.

This discussion points to a fundamental difficulty with the definitional approach. According to this approach, the uniqueness of the subjective probabilities is based on a normalization which is not implied by the structure of the preference relation and, consequently, is devoid of choice-theoretic meaning. For example, the definitions of subjective probabilities in the works of Savage (1954), F. J. Anscombe and R. J. Aumann (1963), Mark Machina and David Schmeidler (1992, 1995), and Itzhak Gilboa and Schmeidler (1989), are based on the convention that constant acts (that is, alternatives that assign the same consequence to every state) yield the same utility in every state. This convention is not testable within the analytical framework of these models.² Nevertheless, insofar as the characterization of the decision maker's choice behavior is concerned, the choice of normalization is immaterial. Why, then, is it important to obtain a "correct" representation of the decision maker's beliefs?

The answer to this question depends on the applications of the theory one has in mind. If one's sole concern is modeling choice behavior in the face of uncertainty,

² In the same vein, Karni's (1993) definition of subjective probability with state-dependent preferences uses the boundedness of the utility function to obtain the required normalization. Karni and Schmeidler (1993) and Robert F. Nau (1995) use the marginal rates of substitution among payoffs in different states to normalize the utility functions.

then the issue of whether the probabilities represent the decision maker's beliefs is irrelevant. In this case, the only significance of the subjective expected utility representation is its additively separable functional form. However, additive separability can be obtained with fewer restrictions on the structure of the underlying preferences. For instance, in the Anscombe and Aumann (1963) framework, additively separable representation characterizes Archimedean weak orders that satisfy the independence axiom of expected utility theory. In particular, state-independence, the axiom requiring that risk attitudes be independent of the state, is unnecessary for the representation. Imposing state-independence renders the model inapplicable for the analysis of important decision problems, such as the demand for health and life insurance, in which the state affects the decision maker's risk attitudes.

Readiness to restrict the theory so as to separate utilities from probabilities stems from the original motivation for developing choice-based subjective probabilities, namely, to furnish choice-based foundations of the Bayesian prior. From this viewpoint, if the probabilities defined by the model are devoid of choice-theoretic meaning, then we are back to square one. Put differently, additively separable representation of the preference relation can be interpreted as expected utility representation with respect to a uniform prior and state-dependent utility functions. Unless the prior probabilities measure the decision maker's beliefs, using the uniform prior seems at least as compelling and, in practice, simpler than using subjective probabilities obtained by an arbitrary normalization of the utility function.

In this paper, I pursue the development of an alternative model based on the measurement approach to the definition of subjective probabilities explored in Karni (2011). In that work, I dispense with Savage's notion of a state space, proposing instead a new analytical framework that consists of a set, Θ , of effects (physical phenomena on which the decision maker may place bets and which may, or may not, impact his well-being); a set, A , of actions (initiatives by which the decision maker believes he can affect the likelihoods of ensuing effects); a set, B , of bets on the effects; and a set of informative and not informative signals, \bar{X} , received before taking actions and choosing bets. The choice set, \mathcal{I} , consists of information-contingent plans (strategies) for choosing actions and bets (that is, a strategy is a function $I : \bar{X} \rightarrow A \times B$). In this framework, Karni (2011) develops a choice-based, Bayesian, decision theory in which decision makers' preferences are represented by

$$(1) \quad I \mapsto \sum_{x \in \bar{X}} \left[\sum_{\theta \in \Theta} \pi(\theta, x | a_{I(x)}) \left[u(b_{I(x)}(\theta), \theta) + v(a_{I(x)}) \right] \right],$$

where $a_{I(x)}$ and $b_{I(x)}$ are the action and bet assigned to the observation x by the strategy I ; $\{u(\cdot, \theta)\}_{\theta \in \Theta}$ are effect-dependent utility functions on the monetary payoffs of the bets; v is the (dis)utility of actions; and $\{\pi(\cdot, \cdot | a)\}_{a \in A}$ is a unique family of action-dependent, joint, subjective probabilities distributions on $\Theta \times \bar{X}$ such that the prior distributions $\{\pi(\cdot | o, a)\}_{a \in A}$ and the posterior distributions $\{\pi(\cdot | x, a)\}_{a \in A}$ on Θ are linked by Bayes rule and represent the decision maker's prior and posterior beliefs. It is worth underscoring the fact that the family, $\{\pi(\cdot, \cdot | a)\}_{a \in A}$, of action-dependent, joint subjective probability distributions is the only family of

such distributions that incorporates new information solely through its effect on the decision-maker's beliefs rather than his tastes.

This paper builds upon and extends the work of Karni (2011). The following example lends concrete meaning to the abstract terms mentioned above and serves to motivate the extensions of this work. Consider a decision maker faced with the prospect of an approaching hurricane. The decision maker must make a plan that, contingent on the weather reports, may include boarding up his house, moving his family to a shelter, and betting on the storm's damage (that is, taking out insurance). The uncertainty is resolved once the weather forecast is obtained, the plan is put into effect, the storm passes, and its path and force have been determined.

In this example, effects correspond to the potential material and bodily damage and actions are the initiatives (e.g., boarding up the house, moving to a shelter) the decision maker can take to mitigate the damage. Bets are alternative insurance policies and observations are weather forecasts. The uncertainty in this example is resolved in two stages. In the first stage a weather forecast obtains, upon the receipt of which, an action and a bet, prescribed by the strategy, are put into effect. In the second stage, the path and force of the hurricane are determined, the associated damage is realized and insurance indemnity is paid.

Consider the issue of subjective probabilities. At the point at which he contemplates his strategies, there are two aspects of uncertainty about which, presumably, the decision maker entertains beliefs. The first concerns the likelihoods of alternative weather reports and, conditional on these reports, the likelihoods of subsequent path-force combinations of the approaching hurricane. The second is the likelihoods of the ensuing levels of damage (the effects). Clearly, the likelihoods of the latter are determined by those of the former coupled with the actions that were taken, in the interim, by the decision maker.

Karni (2011) deals solely with the second aspect of uncertainty and beliefs. The main objective of this paper is to extend the analytical framework to include a state space and define a unique, choice-based, subjective probability measure on the state space generating the family of joint probability distributions $\{\pi(\cdot, \cdot | a)\}_{a \in A}$ that figure in the representation. In terms of the hurricane example, my objective is to infer unique subjective probabilities on the path-force combinations of the approaching hurricane that are consistent with $\{\pi(\cdot, \cdot | a)\}_{a \in A}$, thus resolving the fundamental difficulty with the definition of subjective probabilities mentioned above.

As in Karni (2011), the approach is anchored in the revealed-preference methodology and is Bayesian. The revealed-preference aspect maintains that a decision-maker's choice among alternative courses of action reflects his tastes for the ultimate outcomes and his beliefs regarding the likelihoods of the events in which these payoffs materialize. Consequently, the utility representing the decision maker's tastes and the probabilities representing his beliefs may be inferred from his choice behavior. The Bayesian aspect of the model takes as premises that: (i) new information affects the decision maker's preferences, or choice behavior, through its effect on his beliefs rather than his tastes, and (ii) the posterior probabilities, representing the decision maker's posterior beliefs, are obtained by the updating of the prior probabilities, representing his prior beliefs, using Bayes' rule. It is worth underscoring that the presumption that the decision-maker's posterior preference

relations are obtained from the prior preference relation by the updating of the subjective probabilities leaving the utility functions intact, is critical for the uniqueness of the probabilities.

Section I describes the analytical framework, including the definition of a state-space. It provides a brief summary of the preference structure and representation of Karni (2011). Section II defines the measure space implied by the model and characterizes the unique probability measure on this space that generates the family of action-dependent subjective probability distributions on the effects introduced in Section I. Section III includes concluding remarks. The proofs are collected in the Appendix.

I. The Model

A. The Analytical Framework

Following Karni (2011), let Θ be a finite set of *effects*; let A be a connected separable topological space, whose elements are referred to as *actions*; let X be a finite set of *observations* or *signals*; denote by o the event that no observation materializes and define $\bar{X} = X \cup \{o\}$ ³. A *bet* is a real-valued mapping on Θ , interpreted as monetary payoffs contingent on the realized effect. Let B denote the set of all bets and assume that it is endowed with the $\mathbb{R}^{|\Theta|}$ topology. Denote by $b_{-\theta}r$ the bet obtained from $b \in B$ by replacing the θ -coordinate of b , $b(\theta)$, with r .

Informative and noninformative signals may be received by the decision maker before he chooses a bet and an action, and affect his choice. The decision maker is supposed to formulate a strategy specifying the action-bet pairs to be implemented contingent on the observations. Formally, a *strategy* is a function $I : \bar{X} \rightarrow A \times B$ whose interpretation is a set of instructions specifying, for each observation, x , an action-bet pair, $I(x) = (a_{I(x)}, b_{I(x)})$, to be implemented if the observation x obtains. The set of all strategies, $\mathcal{I} := (A \times B)^{\bar{X}}$, is endowed with the product topology.

A decision maker is characterized by a preference relation \succsim on \mathcal{I} . The strict preference relation, \succ , and the indifference relation, \sim , are the asymmetric and symmetric parts of \succsim , respectively.

As usual, a consequence depicts those aspects of the decision problem that affect the decision maker's ex-post well-being. In this model, a *consequence* is a triplet, (a, r, θ) , representing, respectively, the action, the monetary payoff of the bet, and the effect. The set of all consequences is given by the Cartesian product $C = A \times \mathbb{R} \times \Theta$.

A *state of nature*, or a *state*, is a complete resolution of uncertainty, "a description of the world so complete that, if true and known, the consequences of every action would be known" (Kenneth J. Arrow 1971, 45). In this context, a state, s , is a function from \mathcal{I} to C . The set $S := \{s : \mathcal{I} \rightarrow C\}$ is the *state space*. Subsets of S are *events*. One of the elements of S is the *true state*. An event is said to *obtain* if the true state is an element of it.

³ The interpretation of these terms is as in the introduction.

Uncertainty in this model is resolved in two stages. In the interim stage, an observation, $x \in \bar{X}$, obtains and the action and bet prescribed by each strategy for that observation are determined. In the second stage, the effect is realized and the payoff of the bet is determined. Let Ω be the set of all functions from the set of actions to the set of effects (that is, $\Omega := \{\omega : A \rightarrow \Theta\}$). Elements of Ω depict the resolution of uncertainty surrounding the effects. Thus, $S = \bar{X} \times \Omega$ and each state $s = (x, \omega)$ is an intersection of an *informational event* $\{x\} \times \Omega$ and a *material event* $\bar{X} \times \{\omega\}$. In other words, a state has two distinct dimensions corresponding to the two stages of the resolution of uncertainty, the purely informational dimension, x , and the substantive dimension, ω . Informational events do not affect the decision maker's well-being directly whereas material events might.

In general, states are abstract resolutions of uncertainty. In some situations, however, it is natural to attribute to the states concrete interpretations. In the example of the hurricane in the introduction, the informational events are weather forecasts, and the material events correspond to the path and force of the hurricane. Another example, given by R. Duncan Luce and David H. Krantz (1971), envisions a passenger who, to get from here to there, must choose among driving, taking a bus, or flying. Whether and when he arrives at his destination depends on conditions beyond his control, such as the weather, the mechanical functioning of the alternative means of transportation, roads congestion, and so forth. Conceivably, before choosing the mean of transportation and placing a bet on the outcome (e.g., by taking out insurance) the passenger may get some relevant information (e.g., weather forecast, FAA report on near misses at the destination airport, road construction along his route, etc.) that may affect his decision. The uncertainty regarding the outcome of the trip is resolved once the weather forecast is obtained and the mechanical functioning of the alternative means of transportation, roads congestion, whether there has been a plane crash, and so forth become known.

A main concern of this paper is to define a σ -algebra, \mathcal{E} , on S and a unique probability measure, P , on the measurable space (S, \mathcal{E}) such that (i) the conditioning of P on the noninformative signal o represents the decision maker's prior beliefs and (ii) the conditioning of P on informative signals $x \in X$ represent the decision maker's posterior beliefs.

Denote by $I_{-x}(a, b)$ the strategy in which the x -coordinate of I , $I(x)$, is replaced by (a, b) . The truncated strategy I_{-x} is referred to as substrategy. For every given $x \in \bar{X}$, denote by \succsim^x the induced preference relation on $A \times B$ defined by $(a, b) \succsim^x (a', b')$ if and only if $I_{-x}(a, b) \succsim I_{-x}(a', b')$. The induced strict preference relation, denoted by \succ^x , and the induced indifference relation, denoted by \sim^x , are the asymmetric and symmetric parts of \succsim^x , respectively.⁴ The induced preference relation \succsim^o is referred to as the *prior* preference relation; the preference relations $\succsim^x, x \in X$ are the *posterior* preference relations. An observation, x , is *essential* if $(a, b) \succ^x (a', b')$, for some $(a, b), (a', b') \in A \times B$. I assume throughout that all elements of \bar{X} are essential.

⁴For preference relations satisfying the axioms in Karni (2011), described below, these relations are well-defined. In particular, they are independent of I .

For every $a \in A$ and $x \in \bar{X}$, define a binary relation \succsim_a^x on B by: for all $b, b' \in B$, $b \succsim_a^x b'$ if and only if $(a, b) \succsim^x (a, b')$. The asymmetric and symmetric parts of \succsim_a^x are denoted by \succ_a^x and \sim_a^x , respectively.

An effect, θ , is said to be *nonnull* given the observation-action pair (x, a) if $(b_{-\theta r}) \succ_a^x (b_{-\theta r'})$, for some $b \in B$ and $r, r' \in \mathbb{R}$; it is *null* given the observation-action pair (x, a) otherwise. Given a preference relation, \succsim , denote by $\Theta(a, x)$ the subset of effects that are nonnull given the observation-action pair (x, a) . Assume that $\Theta(a, o) = \Theta$, for all $a \in A$. Note that if $\omega(a) \notin \Theta(a, x)$, then the state (x, ω) is null.

B. Preferences on Strategies and Their Representation

The structure of the preference relations is that of Karni (2011). For the sake of completeness it is described here briefly and informally. Let \succsim on \mathcal{I} be a continuous weak order satisfying coordinate independence (CI), (that is, the requirement that preferences between strategies be independent of the coordinates on which they agree⁵). In addition, let \succsim exhibit independent betting preferences (IBP), requiring that the “intensity of preferences” for monetary payoffs contingent on any given effect be independent of the action, and belief consistency (BC), asserting that, given any action, the preference relation on bets conditional on no new information is the same as that when new information may not be used to select the bet (that is, in and of itself, information is worthless).

Bets whose payoffs offset the direct impact of the effects are dubbed *constant utility bets*. The analytical framework of this paper renders this notion a choice-base phenomenon. To grasp this claim, recall that actions affect decision makers in two ways: directly through their utility cost and indirectly by altering the probabilities of the effects. Moreover, only the indirect impact depends on the observations. In the case of constant utility bets, and only in this case, the intensity of the preferences over the actions is observation-independent. This means that the indirect influence of the actions is neutralized, which can happen only if the utility associated constant utility bets is invariable across the effects. Let $B^{cu}(\succsim)$ be a subset of all constant utility bets according to \succsim . For sufficiently rich set of actions and observations, $B^{cu}(\succsim)$ is well-defined. The set of constant utility bets is said to be *inclusive* if for every $(x, a) \in X \times A$ and $b \in B$ there is $\bar{b} \in B^{cu}(\succsim)$ such that $b \sim_a^x \bar{b}$.

Invoking the notion of constant utility bets, the final two axioms require that the trade-offs between the actions and substrategies be independent of the constant utility bets, and that the direct effect of actions, measured by the preferential difference between constant utility bets, be independent of the observation. These requirements are referred to, respectively, as trade-off independence (TI), and conditional monotonicity (CM).

The main theorem in Karni (2011) asserts that if $B^{cu}(\succsim)$ is inclusive, then \succsim on \mathcal{I} is a continuous weak order satisfying the aforementioned axioms if and only if there exist continuous, real-valued functions $\{u(\cdot, \theta) \mid \theta \in \Theta\}$ on \mathbb{R} , $v \in \mathbb{R}^A$, and a unique

⁵ It is analogous to but weaker than Savage’s (1954) sure-thing principle. See Peter P. Wakker (1989) for details.

family, $\{\pi(\cdot, \cdot | a) | a \in A\}$ of joint probability measures on $\bar{X} \times \Theta$ such that \succsim on I is represented by

$$(2) \quad I \mapsto \sum_{x \in \bar{X}} \mu(x) \left[\sum_{\theta \in \Theta} \pi(\theta | x, a_{I(x)}) u(b_{I(x)}(\theta), \theta) + v(a_{I(x)}) \right],$$

where $\mu(x) = \sum_{\theta \in \Theta} \pi(x, \theta | a)$ for all $x \in \bar{X}$ is independent of a , $\pi(\theta | x, a) = \pi(x, \theta | a) / \mu(x)$ for all $(x, a) \in \bar{X} \times A$, $\pi(\theta | o, a) = (1 / (1 - \mu(o))) \sum_{x \in \bar{X}} \pi(x, \theta | a)$ for all $a \in A$. Moreover, the family of utility functions is unique in the sense that if $\{\hat{u}(\cdot, \theta) | \theta \in \Theta\}$, and $\hat{v} \in \mathbb{R}^A$ is another set of utilities that, together with $\{\pi(\cdot, \cdot | a) | a \in A\}$, represent \succsim , then $\hat{u}(\cdot, \theta) = m\hat{u}(\cdot, \theta) + k, m > 0$ and $\hat{v} = mv + k'$.

Note that, for every constant utility bet, \bar{b} , $u(\bar{b}(\theta), \theta) = u(\bar{b})$, for all $\theta \in \Theta$.

II. Subjective Probabilities on the State Space

A. Action-Dependent Subjective Probabilities on S

The family of joint probability distributions on observations and effects that figure in the representation (2) of the preference relation, can be projected on the underlying state space to obtain a corresponding family of action-dependent, subjective, probability measures. Moreover, this family of measures is the only such family that is consistent with the (unique) joint probability distributions on observations and effects.

To construct the aforementioned family of probability measures, I invoke the fact that the state space can be partitioned to informational events, corresponding to the observations, and, for each action, it can be partitioned into material events corresponding to the effects. Formally, let $\mathcal{Y} = \{\{x\} \times \Omega | x \in \bar{X}\}$ then \mathcal{Y} constitutes an *informational partition* of S . To construct the material partitions, fix $a \in A$ and, for every $\theta \in \Theta$, let $\mathcal{T}_a(\theta) := \{\omega \in \Omega | \omega(a) = \theta\}$. Then $\mathcal{T}_a = \{\bar{X} \times \mathcal{T}_a(\theta) | \theta \in \Theta\}$ is a (finite) *material partition* of S .

Corresponding to each action, we define next a σ -algebra of events. Formally, let \mathcal{E}_a be σ -algebra on S generated by $\mathcal{Y} \wedge \mathcal{T}_a$, the join of \mathcal{Y} and \mathcal{T}_a .⁶ Elements of \mathcal{E}_a are *events*. Hence, events are unions of elements of $\mathcal{Y} \wedge \mathcal{T}_a$.

Consider the measurable spaces (S, \mathcal{E}_a) , $a \in A$. Define a probability measure on \mathcal{E}_a , η_a , by $\eta_a(E) = \sum_{x \in \bar{X}} \sum_{\theta \in \Upsilon} \pi(x, \theta | a)$ for every $E = Z \times \mathcal{T}_a(\Upsilon)$, where $Z \subseteq \bar{X}$, and $\mathcal{T}_a(\Upsilon) = \cup_{\theta \in \Upsilon} \mathcal{T}_a(\theta)$, $\Upsilon \subseteq \Theta$. Then, by representation (2), η_a is unique and the subjective probabilities, $\eta_a(E_I)$, of the informational events $E_I := \{Z \times \Omega | Z \subseteq \bar{X}\}$ are independent of a . We denote these probabilities by $\eta(E_I)$.

For every given a , consider the collection of material events $\mathcal{M}_a := \{\mathcal{T}_a(\Upsilon) | \Upsilon \subseteq \Theta\}$. By representation (2) the prior probability measure on \mathcal{M}_a is given by

$$\eta_a(\mathcal{T}_a(\Upsilon) | o) = \frac{\eta_a(\{o\} \times \mathcal{T}_a(\Upsilon))}{\eta(\{o\} \times \Omega)} = \frac{\sum_{x \in \bar{X}} \sum_{\theta \in \Upsilon} \pi(x, \theta | a)}{1 - \sum_{\theta \in \Theta} \pi(o, \theta | a)} = \sum_{\theta \in \Upsilon} \pi(\theta | o, a).$$

⁶ The join of two partitions is the coarsest common refinement of these partitions.

Similarly, for every $x \in X$, the posterior probability measure on \mathcal{M}_a is given by

$$\eta_a(\mathcal{T}_a(\Upsilon) | x) = \frac{\eta_a(\{x\} \times \mathcal{T}_a(\Upsilon))}{\eta(\{x\} \times \Omega)} = \frac{\sum_{\theta \in \Upsilon} \pi(x, \theta | a)}{\sum_{\theta \in \Theta} \pi(x, \theta | a)} = \sum_{\theta \in \Upsilon} \pi(\theta | x, a).$$

Then the main representation theorem in Karni (2011) may be restated in terms of these probability measures as follows:⁷

THEOREM 1: *Let \succsim be a preference relation on I and suppose that $B^{cu}(\succsim)$ is inclusive. Then \succsim is a continuous weak order satisfying CI, IBP, BC, TI and CM if and only if there exist continuous, real-valued functions u on $\mathbb{R} \times \Theta$ and v on A , and for each $a \in A$, a unique probability measure η_a on the measurable space (S, \mathcal{E}_a) such that $\eta_a(\{x\} \times \Omega) = \eta_{a'}(\{x\} \times \Omega) = \eta(x)$, for all $a, a' \in A$ and $x \in \bar{X}$, and \succsim on \mathcal{I} is represented by*

$$(3) \quad I \mapsto \sum_{x \in \bar{X}} \eta(x) \left[\sum_{\theta \in \Theta} u(b_{I(x)}(\theta), \theta) \eta_{a_{I(x)}}(\mathcal{T}_{a_{I(x)}}(\theta) | x) + v(a_{I(x)}) \right].$$

Moreover, if \hat{u} on $\mathbb{R} \times \Theta$ and \hat{v} on A , is another pair of real-valued functions that, together with η_a , represent \succsim , then $\hat{u}(\cdot, \theta) = m\hat{u}(\cdot, \theta) + k$, $m > 0$ and $\hat{v} = mv + k'$. u is unique up to positive affine transformation.

Theorem 1 asserts the existence of a unique collection of measure spaces $\{ (S, \mathcal{E}_a, \eta_a) | a \in A \}$. This collection is sufficiently rich to allow action-dependent probability to be defined for every event that matters to the decision maker, for all the conceivable choices among strategies that he might be called upon to make. Hence, from the viewpoint of Bayesian decision theory, the family of action-dependent subjective probability measures defined here is complete in the sense of being well defined for every conceivable decision problem that can be formulated in this framework. However, there is no guarantee that these subjective probability measures are mutually consistent. Hence, it is interesting to inquire about the necessary and sufficient conditions for the existence of a unique probability space that supports all these action-dependent measures in the sense that $\eta_a(E)$ coincides with this measure for every $a \in A$ and $E \in \mathcal{E}_a$. This issue is taken up next.

B. Subjective Probabilities on S

To begin with, it is necessary to define a σ -algebra that includes the class of events $\cup_{a \in A} \mathcal{E}_a$. Let $\wedge_{i=1}^k \mathcal{T}_{a_i}$ be the join of $\mathcal{T}_{a_1}, \dots, \mathcal{T}_{a_k}$, and denote by \mathcal{E} the σ -algebra generated by $\{ \wedge_{i=1}^k \mathcal{T}_{a_i} | k < \infty \} \wedge \mathcal{Y}$. The issue posed at the end of the last subsection may be restated formally as follows: What are the necessary and sufficient conditions for the existence of a unique probability measure P on the measurable space (S, \mathcal{E}) , such that for all $a \in A$, if $E \in \mathcal{E}_a$, then $\eta_a(E) = P(E)$ and, for all $x \in \bar{X}$, $P(\{x\} \times \Omega) = \eta(x)$?

⁷The proof follows from Theorem 3 in Karni (2011b) by simple substitutions and is omitted.

This question raises two issues. The first, mentioned above, concerns the consistency of beliefs across actions. To illustrate this issue and motivate the ensuing inquiry, I consider the example of the approaching hurricane described in the introduction. In this example the states are concrete, namely, the weather forecast coupled with the force and path of the approaching hurricane. The effects of the storm are determined by the action taken, a , and the force and path of the hurricane, ω . In this case, for each $(\theta, x, a) \in \Theta \times X \times A$, the probability $\pi(\theta|x, a)$, corresponds to probability $\eta_a(\mathcal{T}_a(\theta)|x)$. Moreover, because the force and trajectory of the hurricane are independent of the action, and since the probabilities of the effects are induced by the probabilities of the underlying state, consistency requires that if $\mathcal{T}_a(\theta) = \mathcal{T}_{a'}(\theta') = T \in \mathcal{M}_a \cap \mathcal{M}_{a'}$ then $\pi(\theta|x, a) = \eta_a(T|x) = \eta_{a'}(T|x) = \pi(\theta'|x, a')$.

The second issue arises because a probability measure on the measurable space (S, \mathcal{E}) requires that the probabilities of events such as $\mathcal{T}_a(\theta) \cap \mathcal{T}_{a'}(\theta') \neq \emptyset$, for distinct actions, $a, a' \in A$, be well-defined. However, the actions a and a' are mutually exclusive. Therefore, to assign unique probability to material events that are in the intersection of the σ -algebras corresponding to distinct actions, it must be the case that the collection $\cup_{a \in A} \mathcal{E}_a$ is rich in the sense the same material event, namely, $\mathcal{T}_a(\theta) \cap \mathcal{T}_{a'}(\theta')$, is an element of the σ -algebra, say $\mathcal{E}_{a''}$, corresponding to some action a'' in A .

Action-independent Beliefs.—To formalize the notion of belief consistency I consider the situation in which, upon receiving a signal, x , a decision maker takes the action a and places a bet on the material event $T \in \mathcal{M}_a$. Formally, given $\bar{b}, \bar{b}' \in B^{cu}(\succsim)$ and $T \in \mathcal{M}_a$, define $\bar{b}_T \bar{b}' \in B$ as follows:

$$(\bar{b}_T \bar{b}')(\theta) = \left\{ \begin{array}{ll} \bar{b}(\theta) & \text{if } \mathcal{T}_a(\theta) \subseteq T \\ \bar{b}'(\theta) & \text{if } \mathcal{T}_a(\theta) \subset \Omega - T \end{array} \right\}.$$

Thus, conditional on x and a , $\bar{b}_T \bar{b}'$ is the bet whose utility payoffs is a doubleton, paying $u(\bar{b}(\theta))$ for all θ such that $\mathcal{T}_a(\theta) \subseteq T$ and $u(\bar{b}'(\theta))$, for all θ such that $\mathcal{T}_a(\theta) \subset \Omega - T$. If $I_{-x}(a, \bar{b}) \succ I_{-x}(a, \bar{b}')$ then $\bar{b}_T \bar{b}'$ is a bet on the material event T . For all $x \in \bar{X}$ and $T \in \mathcal{M}_a \cap \mathcal{M}_{a'}$ let $E(x, T) := \{x\} \times T$.

DEFINITION 2: *The event $E(x, T) \in \mathcal{E}_a \cap \mathcal{E}_{a'}$ is believed equally likely to obtain under a and under a' if $I_{-x}(a, \bar{b}_T \bar{b}') \sim I_{-x}(a', \bar{b}_T \bar{b}')$, where $I_{-x}(a, \bar{b}) \sim I_{-x}(a', \bar{b})$, $I_{-x}(a, \bar{b}') \sim I_{-x}(a', \bar{b}')$ and $I_{-x}(a, \bar{b}) \succ I_{-x}(a, \bar{b}')$.*

Let the (partial) binary relation on $\cup_{a \in A} \mathcal{E}_a$ given by Definition 2 be denoted by $(a, E(x, T)) \sim_L (a', E(x, T))$. Then \sim_L is an equivalence relation. The next axiom requires that if the same event may obtain under two distinct actions, then the beliefs about its likelihood be action-independent. Formally,

Action-independent Beliefs.—For all $a, a' \in A$ and $E(x, T) \in \mathcal{E}_a \cap \mathcal{E}_{a'}$, $(a, E(x, T)) \sim_L (a', E(x, T))$.

The following theorem asserts that, under action-independent beliefs, the action-dependent subjective probability measures, $\{\eta_a\}_{a \in A}$, agree on events in the intersections of the σ -algebras $\{\mathcal{E}_a\}_{a \in A}$.

THEOREM 3: *Let \succsim be a continuous weak order on \mathcal{I} satisfying CI, IBP, BC, TI, and CM and suppose that $B^{cu}(\succsim)$ is inclusive. Then \succsim satisfies action-independent beliefs if and only if $\eta_a(E(x, T)) = \eta_{a'}(E(x, T))$, for all $a, a' \in A$ and $E(x, T) \in \mathcal{E}_a \cap \mathcal{E}_{a'}$.*

Richness.—Consider next the problem posed by events in \mathcal{E} that are in the intersection of elements of the partitions of the state space under distinct actions. In general, the probabilities of such events cannot be inferred from the action-dependent joint probabilities of observations and effects that figure in the representation (2). This problem can be surmounted if the decision model, (S, A, Θ, \bar{X}) , consisting of a state space, S , action set, A , a finite set of effect, Θ , and a finite set of observations, \bar{X} , is rich in the sense that, for every finite collection of actions and corresponding subsets of effects, the event that yields these effects under these actions coincides with an event in \mathcal{E}_a for some $a \in A$. Formally,

DEFINITION 4: *The model (S, A, Θ, \bar{X}) is rich if, for all $(a_i, \Upsilon_i) \in A \times 2^\Theta$, $i = 1, \dots, k$, $k < \infty$, and $x \in \bar{X}$, $\{x\} \times \cap_{i=1}^k \mathcal{T}_{a_i}(\Upsilon_i) = E$ for some $E \in \cup_{a \in A} \mathcal{E}_a$.*

The next theorem establishes the existence and uniqueness of a probability space supporting the action-dependent distributions.

THEOREM 5: *Let \succsim be a preference relation on \mathcal{I} and suppose that $B^{cu}(\succsim)$ is inclusive and that the model (S, A, Θ, \bar{X}) is rich. Then \succsim is a continuous weak order satisfying CI, IBP, BC, TI, and CM and action-independent beliefs if and only if there exist continuous, real-valued functions u on $\mathbb{R} \times \Theta$ and v on A , and a unique probability measure P on (S, \mathcal{E}) such that \succsim on \mathcal{I} is represented by*

$$(4) \quad I \mapsto \sum_{x \in \bar{X}} P(x) \sum_{\theta \in \Theta} \left[\sum_{\theta \in \Theta} u(b_{I(x)}(\theta), \theta) P(\mathcal{T}_{a_{I(x)}}(\theta) | x) + v(a_{I(x)}) \right],$$

where $P(x) := P(\{x\} \times \Omega)$. Moreover, if \hat{u} on $\mathbb{R} \times \Theta$ and \hat{v} on A , is another pair of real-valued functions that, together with η_a , represent \succsim , then $\hat{u}(\cdot, \theta) = m\hat{u}(\cdot, \theta) + k, m > 0$ and $\hat{v} = mv + k'$.

The proof invokes the probability distributions the figure in the representation (2) and the richness of the model to define a probability measure on the algebra generated by $\{\wedge_{i=1}^k \mathcal{T}_{a_i} | k < \infty\} \wedge \mathcal{Y}$. This probability measure has a unique extension to (S, \mathcal{E}) .

III. Concluding Remarks

Extending the analysis in Karni (2011), this paper presents a model of Bayesian decision making and choice-based foundations of prior and posterior subjective

probabilities on a state-space. The model admits effect-dependent utility functions of wealth, making it applicable for the analysis of life and health insurance problems, medical decision making, and other decisions in which the effect interacts with the decision maker’s risk attitudes.

The theory of contracts in the presence of moral hazard admits three distinct formulations.⁸ According to the state-space formulation of the principal-agent problem, the outcome, $\theta \in \Theta$, of a contractual agreement is determined by the action of the agent, $a \in A$, and by the state of nature, $\omega \in \Omega$, according to a technology, $t(a, \omega) = \omega(a) = \theta$. The principal’s problem is to design a contract that awards the agent according to the observed outcome, so as to induce the agent to implement the (unobserved) action desired by the principal. Actions are costly in terms of the agent’s well-being. The model of this paper furnishes axiomatic foundations of the agent’s behavior in the state-space formulation of the principal-agent problem.

The definition of unique, choice-based, subjective probabilities on a state-space settles a long standing issue in Bayesian theory.

APPENDIX

PROOF OF THEOREM 3:

Suppose that \succsim is a continuous weak order on \mathcal{I} satisfying CI, IBP, BC, TI, and CM, $B^{cu}(\succsim)$ is inclusive and $E(x, T) \in \mathcal{E}_a \cap \mathcal{E}_{a'}$. Let $\bar{b}, \bar{b}', \bar{b}'' \in B^{cu}(\succsim)$ and $I, I' \in \mathcal{I}$ satisfy the following conditions: (a) $I_{-x}(a, \bar{b}) \sim I'_{-x}(a', \bar{b})$, (b) $I_{-x}(a, \bar{b}'') \sim I'_{-x}(a', \bar{b}'')$, and (c) $I_{-x}(a, \bar{b}'') \succ I_{-x}(a, \bar{b}) \succ I_{-x}(a, \bar{b}')$. Then, by the representation (2),

$$(5) \quad u(\bar{b}'') - u(\bar{b}) > 0.$$

By Definition 2, $(a, E(x, T)) \sim_L (a', E(x, T))$ if and only if $I_{-x}(a, \bar{b}_T^* \bar{b}') \sim I'_{-x}(a', \bar{b}_T^* \bar{b}')$, where $I_{-x}(a, \bar{b}^*) \sim I'_{-x}(a', \bar{b}^*)$, $I_{-x}(a, \bar{b}') \sim I'_{-x}(a', \bar{b}')$ and $I_{-x}(a, \bar{b}^*) \succ I'_{-x}(a, \bar{b}')$, for $\bar{b}^* \in \{\bar{b}, \bar{b}''\}$. For every strategy I and observation x , let

$$k(I_{-x}) = \sum_{x' \in \bar{X} - \{x'\}} \eta(x') \left[\sum_{\theta \in \Theta} u(b_{I(x')}(\theta), \theta) \eta_{a_I}(x') (\mathcal{T}_{a_I(x')}(\theta) | x) + v(a_{I(x')}) \right]$$

By Theorem 1, and conditions (a) and (b),

$$(6) \quad I_{-x}(a, \bar{b}_T \bar{b}') \sim I'_{-x}(a', \bar{b}_T \bar{b}')$$

if and only if

$$(7) \quad \begin{aligned} k(I_{-x}) + \eta_a(E(x, T))u(\bar{b}) + (1 - \eta_a(E(x, T)))u(\bar{b}') + v(a) \\ = k(I'_{-x}) + \eta_{a'}(E(x, T))u(\bar{b}) + (1 - \eta_{a'}(E(x, T)))u(\bar{b}') + v(a') \end{aligned}$$

⁸ See Oliver Hart and Bengt Holmström (1987) for details and references.

and

$$(8) \quad I_{-x}(a, \bar{b}'' \bar{b}') \sim I'_{-x}(a', \bar{b}'' \bar{b}')$$

if and only if

$$(9) \quad k(I_{-x}) + \eta_a(E(x, T))u(\bar{b}'') + (1 - \eta_a(E(x, T)))u(\bar{b}') + v(a) \\ = k(I'_{-x}) + \eta_{a'}(E(x, T))\tilde{u}(\bar{b}'') + (1 - \eta_{a'}(E(x, T)))u(\bar{b}') + v(a').$$

Hence,

$$(10) \quad \eta_a(E(x, T)) [u(\bar{b}) - u(\bar{b}'')] = \eta_{a'}(E(x, T)) [u(\bar{b}) - u(\bar{b}')].$$

Equation (5) implies that equation (10) holds if and only if $\eta_a(E(x, T)) = \eta_{a'}(E(x, T))$. Thus $(a, E(x, T)) \sim_L (a', E(x, T))$, for all $a, a' \in A$ and $E(x, T) \in \mathcal{E}_a \cap \mathcal{E}_{a'}$, if and only if $\eta_a(E(x, T)) = \eta_{a'}(E(x, T))$. Hence \succsim on \mathcal{I} satisfies action-independent beliefs if and only if $\eta_a(E(x, T)) = \eta_{a'}(E(x, T))$.

PROOF OF THEOREM 5:

Let \mathcal{A} denote the algebra generated by $(\bigwedge_{i=1}^k \mathcal{T}_{a_i}) \wedge \mathcal{Y}$, the join of the partitions $\mathcal{T}_{a_1}, \dots, \mathcal{T}_{a_k}$, where k is an integer, and \mathcal{Y} . Elements of \mathcal{A} include all the intersections of sets belonging to partitions $\mathcal{T}_{a_1}, \dots, \mathcal{T}_{a_k}$. Because (S, A, Θ, \bar{X}) is rich, each event involving such interactions is equal to a “simple” event, $E \in \mathcal{E}_a$, for some $a \in A$. By Theorem 1, the probability of E is unique and is given by $\eta_a(E)$.

Define the probability measure P_0 on \mathcal{A} as follows: for all $E \in \mathcal{A}$,

$$(11) \quad P_0(E) = \eta_a(E).$$

Since \mathcal{E} is the σ -algebra generated by \mathcal{A} , Theorem 3.1 in Partick Billingsley (1986) implies that P_0 has a unique extension, P , to \mathcal{E} . Then P is a probability measure on (S, \mathcal{E}) . In particular, let $P(\mathcal{T}_a(\theta) | x) = P(\{x\} \times \mathcal{T}_a(\theta)) / P(\{x\} \times \Omega)$, then

$$(12) \quad P(\mathcal{T}_a(\theta) | x) = \eta_a(\mathcal{T}_a(\theta) | x) = \pi(x, \theta | a)$$

and

$$(13) \quad P(x) := P(\{x\} \times \Omega) = \eta(x) = \sum_{\theta \in \Theta} \pi(x, \theta | a) = \mu(x).$$

Substitute $P(\mathcal{T}_a(\theta) | x)$ for $\eta_{a_{\pi(x)}}(\mathcal{T}_{a_{\pi(x)}}(\theta) | x)$ and $P(x)$ for $\eta(x)$ in Theorem 1, to obtain the representation (4).

REFERENCES

- Anscombe, F. J., and R. J. Aumann.** 1963. "A Definition of Subjective Probability." *Annals of Mathematical Statistics*, 34(1): 199–205.
- Arrow, Kenneth J.** 1971. *Essays in the Theory of Risk Bearing*. Chicago: Markham Publishing Co.
- Billingsley, Patrick.** 1986. *Probability and Measure*. New York: John Wiley and Sons.
- de Finetti, Bruno.** 1937. "La prévision: ses lois logiques, ses sources subjectives." *Annales de l'Institut Henri Poincaré*, 7(1): 1–68. Translated by Howard E. Kyburg in *Studies in Subjective Probability*. (New York: John Wiley and Sons, 1964).
- Gilboa, Itzhak, and David Schmeidler.** 1989. "Maxmin Expected Utility with Non-unique Prior." *Journal of Mathematical Economics*, 18(2): 141–53.
- Hart, Oliver, and Bengt Holmström.** 1987. "The Theory of Contracts." In *Advances in Economic Theory: Fifth World Congress*, ed. Truman F. Bewley, 71–155. Cambridge, UK: Cambridge University Press.
- Karni, Edi.** 2011a. "Subjective Probabilities on a State Space: Dataset." *American Economic Journal: Microeconomics*. <http://www.aeaweb.org/articles.php?doi=10.1257/mic.3.4.172>.
- Karni, Edi.** 2011b. "A Theory Of Bayesian Decision Making with Action-Dependent Subjective Probabilities." *Economic Theory*, 48(1): 125–46.
- Karni, Edi.** 1993. "A Definition of Subjective Probabilities with State-Dependent Preferences." *Econometrica*, 61(1): 187–98.
- Karni, Edi, and David Schmeidler.** 1993. "On the Uniqueness of Subjective Probabilities." *Economic Theory*, 3(2): 267–77.
- Luce, R. Duncan, and David H. Krantz.** 1971. "Conditional Expected Utility." *Econometrica*, 39(2): 253–71.
- Machina, Mark J., and David Schmeidler.** 1992. "A More Robust Definition of Subjective Probability." *Econometrica*, 60(4): 745–80.
- Machina, Mark J., and David Schmeidler.** 1995. "Bayes without Bernoulli: Simple Conditions for Probabilistically Sophisticated Choice." *Journal of Economic Theory*, 67(1): 106–28.
- Nau, Robert F.** 1995. "Coherent Decision Analysis with Inseparable Probabilities and Utilities." *Journal of Risk and Uncertainty*, 10(1): 71–91.
- Ramsey, Frank P.** 1931. "Truth and Probability." In *The Foundations of Mathematics and Other Logical Essays*, ed. Richard B. Braithwaite, 156–98. London: Kegan, Paul, Trench, Trubner and Co.
- Savage, Leonard J.** 1954. *The Foundations of Statistics*. New York: John Wiley and Sons.
- Wakker, Peter P.** 1989. *Additive Representations of Preferences: A New Foundation of Decision Analysis*. Dordrecht: Kluwer Academic Publishers.