



A theory-based decision model

Edi Karni ¹

Johns Hopkins University, Department of Economics, United States of America

Received 25 June 2021; final version received 3 March 2022; accepted 6 March 2022

Available online 16 March 2022

Abstract

This paper proposes a theory-based model of decision-making under uncertainty the main premise of which is that predictions of the outcomes of acts are derived from theories. Realized act-outcome pairs provide information on the basis of which decision makers update their beliefs regarding the validity of the underlying theories. Consequently, acts are, simultaneously, information-generating initiatives, or experiments, that have material consequences. Experiments, that is, information-generating initiatives of no direct material consequences, are characterized and the value of information they generate defined. An incentive-compatible mechanism is introduced by which the beliefs decision-makers holds regarding the validity of the theories are elicited.

© 2022 Elsevier Inc. All rights reserved.

JEL classification: D8; D81; D83

Keywords: Theory-based decisions; Experimentation; Value of information; Subjective probabilities; Probability elicitation

1. Introduction

The objective of this paper is twofold. It develops a theory-based decision model and analyzes its implications for individual choice under uncertainty, and it develops a theory of choice of experiments. Two traits of human nature are pertinent for this purpose. The first is the desire to explore, by means of observation and experimentation, the physical, biological, and social

E-mail address: karni@jhu.edu.

¹ I am grateful to Itzhak Gilboa, Marzena Rostek, an associate editor, and two anonymous referees for their useful comments and suggestions.

environments and to distill – in the form of general laws – regularities in empirical observations. The second, is to invoke these laws when predicting the consequences of alternative courses of action.

In the context of decision making under uncertainty, theories – a term I use interchangeably with hypotheses or models – are general laws that decision-makers invoke, explicitly or implicitly, to make sense of alternative actions and predict their outcomes. Theories of decision making under uncertainty based on the analytical framework of Savage (1954) presume that decision-makers predict the results of their actions by assigning consequences to events (i.e., subsets of the state space) and attributing subjective probabilities to these events. In these theories, decision-makers' choice behavior is governed by an *unspecified* combination of prophetic and technological predictions.²

The main premise of this paper is that, in many situations involving decision making under uncertainty, decision-makers invoke *formal theories*, or *statistical models*, to predict the outcomes of alternative actions and choose among feasible actions on the basis of these predictions.³ In situations that admit competing theories that yield distinct predictions, decision-makers weight these predictions by their beliefs in the validity of the underlying theories. In these contexts, observations, consisting of action-outcome pairs, are used to update decision-makers' beliefs in the validity of the underlying theories. A theory is falsified when an action results in an outcome that is outside the set of outcomes that, according to the theory, are feasible given the action. Observations that are consistent with the existing theories induce the updating of the beliefs of the decision-makers in the truth of the theories.

It is common to distinguish acts from experiments. Acts are initiatives that have material consequences and no informational content; experiments are information-generating initiatives with no direct material implications. This dichotomy is an idealized simplification. In reality, acts result in material consequences that also inform decision-makers about the validity of the theories and are, therefore, a kind of experiment. In addition to their informative value, experiments may have, material implications (e.g., cost) and are, therefore, a kind of acts.

In this paper, I develop a general theory of choice that amalgamates the material and informational aspects of acts and considers the choice of experiments as a special case of this theory. In the proposed model, the choice of acts may involve a trade-off between exploitation (i.e., taking the action that yields the best material payoff) and exploration (i.e., trying actions to improve the understanding of their uncertain consequences).

In view of the exploitation-exploration features the proposed model is a contribution to the literature on multi-armed bandits with dependent arms. In particular, actions correspond to the bandit's arms, and the decision problem requires sequential choices from a finite set of possible actions. The choice of actions result in material outcomes that constitute the decision maker's reward and an observations, consisting of action-outcome pairs, that provide valuable information relevant to the choice of subsequent actions. In the model of this paper, observations provide information about the validity of the underlying theories which, in turn, induce stochastic dependency of the arms and correlation among the distributions of the rewards.⁴ This feature of

² Following Popper (2002), prophetic predictions refers to forecasting an event one can do nothing to prevent (e.g., the coming of a typhoon); technological predictions intimate the steps one may take to achieve certain outcomes.

³ Another possible interpretation replaces theories by experts, or consultants, who use different models to predict the outcomes of the actions..

⁴ This feature of the model distinguishes it from the standard bandit problem in which “the distribution of returns from one arm only changes when this arm is chosen,” (Bergemann and Välimäki (2008)).

the model is shared by multi-armed bandit models as the global multi-armed bandit model (Atan et al. (2015)) and the multi-armed bandit model in which distribution dependencies among the arms are assumed to be correlated (Pandey et al. (2007)). In these models, there are unknown parameters that determine the distributions of the rewards of the arms. Because of the dependency on the unknown parameters, pulling an arm and observing the result inform the decision maker about the distributions of the rewards corresponding the other arms. The model of this paper is distinguished from those that were proposed in the literature on multi-armed bandits with dependent arms. Unlike the aforementioned models, it inquires into the source of the dependency among the arms rather treating it as primitive. It also derives the objective function from the underlying preference relation and admits subjective evaluations of the payoffs that is not necessarily separately additive. The more general formulation allows the payoffs in one period to affect the evaluation of the payoffs in subsequent periods. For example, a big monetary gain in one period may affect the willingness to take financial risks when it is time to choose the next action. Moreover, the model weights the alternative theories by the belief of the decision maker in their validity, and the belief is represented by subjective probabilities that are derived from the preference relation rather than being assumed ad hoc.⁵

Because the relevant theories are determined by the context (unlike in Savage's grand-world vision of decision making under uncertainty), the proposed decision model is context-dependent. In the context of medical decisions, for instance, the relevant theories predict the probable states of health (i.e., outcomes) that would result from alternative treatments (i.e., acts). In the context of financial decisions, the relevant theories are models of financial markets that predict the probable values (i.e., outcomes) of different portfolio positions (i.e., acts).

To grasp the ideas to be explored, consider the following simple example. An urn is known to contain 100 balls, 50 of which are red and 50 are black. Balls are drawn sequentially, at random, and their colors observed. A decision-maker can place a bet on the event A , (e.g., all the balls are of the same color) or on any other color combination. The decision-maker entertains two hypotheses regarding the random process generating the outcomes. Hypothesis I is that the draws are with replacement; hypothesis II is that the draws are without replacement. The two hypotheses imply distinct distributions on the various events. If, for example, two balls are drawn, then according to hypothesis I , the event A and its complement, A^c , are equally probable. According to hypothesis II , the probability of event A is $49/99$ and that of A^c is $50/99$. In this scenario, the stochastic process takes place behind a "veil of ignorance" and decision-makers get to observe only the colors of the balls. Acts are sequences of random draws of balls, and outcomes are the resulting color combinations.

A focal issue in the theory of decision making under uncertainty is the existence and elicitation of decision-makers' subjective probabilities of events, such as A and A^c . In this example, these probabilities are induced by underlying (prior) beliefs regarding the validity of the two hypotheses.

Suppose that these beliefs are quantifiable by probabilities and let p and $1 - p$ denote the probabilities of hypotheses I and II , respectively. Consider betting on the color combinations of the balls. Following Borel (1924) and Ramsey (1926), the acceptable odds in such bets allow an observer to infer a decision-maker's beliefs regarding the likely realization of events, such as A and A^c , and quantify those beliefs by probabilities. The question I seek to answer here is

⁵ Much of the literature on multi-armed bandits invokes minimized regret or maximizing the discounted expected value as the criteria for ranking the arm-pulling strategies. In neither case is the criterion derived from underlying preferences.

whether an observer can recover from the betting odds the probabilities the decision-maker assigns the underlying hypotheses. For example, suppose that implementing one of the available elicitation schemes (e.g., Karni (2009)), the observer concludes that the decision-maker's subjective probability of the event A is 0.496. The observer can recover the probabilities p by solving the equation $[99p + (1 - p)98] / 198 = 0.496$.⁶

The following examples set the stage and buttress the argument in favor of the proposed model.

COVID-19: The decision problem facing the policy makers in charge of combating a new variant of COVID-19 has the features of the proposed model. The novelty and the lack of familiarity with the way it spreads give rise to alternative hypotheses predicting the evolution and dynamic of the pandemic. These hypotheses incorporate factors such as the transmission modes (e.g., aerosol and/or fomite) the infectious period (e.g., how long a person is contagious before showing symptoms) the effectiveness of the existing vaccines in preventing infection and the need for hospitalization. Implementation of a policy, such as restrictions on gathering in closed places, has immediate economic and health care consequences. Reflecting the diversity of policy makers' beliefs (and other political considerations), different countries adopted distinct policies, ranging from complete lockdown to taking no measures at all. Evidence regarding the efficacy of the different policies is used to update the beliefs in the validity of the alternative models and revise the policies.

Education signals: An employer who regularly hires employees is interested in their productivity. Productivity is idiosyncratic and cannot be ascertained except by actually employing a worker. Suppose, for simplicity, that the employer entertains alternative hypotheses regarding the relationship between employees' productivity and their education levels. Distinct hypotheses maintain that the employees' education levels and productivity are positively correlated but to different degrees. The employer holds a prior belief about the validity of the alternative hypotheses, based on which she decides to require a certain level of education to fill job vacancies. After having observed the productivity of the employees, the employer updates her beliefs about the validity of the underlying hypotheses, which she then relies upon the next time she looks to fill a job vacancy.

Medical decisions: A patient showing certain symptoms seeks medical treatment. The attending physician may entertain different hypotheses regarding the underlying cause, which she holds with different degrees of confidence. Each hypothesis predicts distinct probable outcomes of the available treatments. Once a treatment is administered, its outcome provides information regarding the possible underlying illness and may be used to update the physician's belief about the likely cause of the symptoms.

Monetary policy: To reduce the unemployment rate, the central bank considers quantitative easing to inflate prices. Competing theories regarding the effect of such policies on unemployment range from the Phillips curve model, which predicts a persistent negative correlation between the rate of inflation and the rate of unemployment, to the long-run Phillips curve model, which predicts that a higher rate of inflation may reduce the rate of unemployment temporarily but that, once inflationary expectations are formed, the rate of unemployment reaches its natural level at the higher inflation rate. The monetary authority entertains probabilistic beliefs about the

⁶ Applying Bayes' rule to obtain the posterior beliefs regarding the validity of hypotheses I and II , it is possible to predict the posterior probabilities that a Bayesian decision maker assigns to events A and A^c and the odds he will accept to bet on these events. Eliciting the posterior probabilities of these events and comparing them to the predictions is a test of Bayesianism. Section 5 provides a more general treatment of the elicitation issue.

validity of these models and, based on those beliefs, implements a policy. Once the effects of the policy are observed and analyzed, the central bank updates its beliefs regarding the validity of the alternative models and invokes its posterior beliefs the next time it is called upon to intervene.

In these examples, the theoretical or statistical models do not necessarily assign the acts unique outcomes. Rather, they predict a distribution on a set of possible outcomes the randomness of which be from factors that are either unobserved or have not been properly accounted for by these models.

The rest of the paper is organized as follows. The next section describes the analytical framework. Section 3 depicts the properties of the preference relations and their representations. Section 4 models experiments and discusses the value of information. Section 5 introduces a novel, incentive-compatible mechanism designed to elicit the decision-maker’s subjective degrees of beliefs in the truth of the relevant theories. Section 6 includes further discussion of the model and a review the related literature. Section 7 provides the proofs.

2. The analytical framework

2.1. Theories, observations, and decision models

Theories are laws depicting the causal relationships between acts and outcomes on the basis of which decision-makers predict the consequences of their actions. In general, theoretical predictions include a stochastic component, reflecting the fact that factors not accounted for by a theory may play a role in determining the action’s outcomes. To formalize this idea, I introduce three primitives: a finite set, $T = \{t_1, \dots, t_m\}$, of *theories*; a finite set, $X = \{x_1, \dots, x_n\}$, of *outcomes*; and a set F of *acts*, whose elements are mappings from the set of theories to the set of Borel-measurable functions on the unit interval, Ω , taking values in X . Formally, for each $f \in F$ and $t \in T$, let $f(t)(x) := \mu_t(f^{-1}(x))$, for all $x \in X$, where μ_t is a Borel probability measure on Ω , representing the probability that theory t assigns to the event $f^{-1}(x)$. Under this definition, F is identified with $(\Delta X)^T$, where ΔX denotes the simplex in \mathbb{R}^n .

For all $f, f' \in F$, and $\alpha \in [0, 1]$, define $(\alpha f + (1 - \alpha) f') \in F$, by:

$$(\alpha f + (1 - \alpha) f')(t)(x) = \alpha \mu_t(f^{-1}(x)) + (1 - \alpha) \mu_t(f'^{-1}(x)), \forall (t, x) \in T \times X.$$

Thus, F is a convex set in $\mathbb{R}^{|T \times X|}$.

Observations are act-outcome pairs, $(f, x) \in F \times X$. For every $f \in F$ and $t \in T$, the support of $f(t)$ is the set $X(t, f) := \{x \in X \mid \mu_t(f^{-1}(x)) > 0\}$. An observation (f, x) is *consistent* with theory t if $x \in X(t, f)$. If (f, x) is inconsistent with theory t , (that is, $x \notin X(t, f)$), then theory t is said to be *falsified* by the observation (f, x) .

A *decision model* is a triplet $\{F, X, T\}$ whose acts, outcome and theories are context specific. Decision models encompasses two layers of randomness. The first layer, modeled by $(\mu_t \mid t \in T)$, represents the *objective* randomness inherent in the predictions of the theories. The second layer is the decision-maker’s *subjective* uncertainty regarding the truth of the theories.

2.2. The choice set and preference relations

Decisions are choices of finite sequences of acts. To simplify the exposition, without essential loss of generality, I model decisions as two-stage dynamic processes.⁷ In the first stage, the decision-maker chooses an act, $f \in F$, and obtains an outcome, $x \in X$. In the second stage, the decision-maker chooses the subsequent act contingent on the observations. Formally, let $\mathcal{Z} := \{\zeta : F \times X \rightarrow F\}$ the set of mappings representing choices of the second-stage acts contingent on the first-stage observations. For each $f \in F$, let $\mathcal{Z}(f)$ denote the set of contingent plans based on the observations that are obtainable once the first-stage act is chosen. Formally, $\mathcal{Z}(f) := \{\zeta(f) : X \rightarrow F \mid \zeta \in \mathcal{Z}\}$. For every given f , the constant contingent plan $\zeta(f, x) = \zeta(f, x')$, for all $x, x' \in X$ is denoted by $\bar{\zeta}(f) \in F$.

The choice set is $\mathbb{C} := \{(f, \zeta(f)) \in F \times \mathcal{Z}(f)\}$, the generic element of which consists of an act f and a plan for choosing a second-stage act $\zeta(f, x)$ contingent on the outcome $x \in X$. For all $(f, \zeta(f)), (f', \zeta'(f')) \in \mathbb{C}$ and $\alpha \in [0, 1]$, define the convex operation

$$\alpha(f, \zeta(f)) + (1 - \alpha)(f', \zeta'(f')) = (\alpha f + (1 - \alpha)f', \alpha\zeta(f) + (1 - \alpha)\zeta'(f')),$$

where $\alpha\zeta(f) + (1 - \alpha)\zeta'(f') := (\alpha\zeta(f, x) + (1 - \alpha)\zeta'(f', x))_{x \in X}$. Then \mathbb{C} is a convex set in a linear space.

A preference relation \succsim on \mathbb{C} is a binary relation that has the following interpretation. For all $(f, \zeta(f)), (f', \zeta'(f')) \in \mathbb{C}$, $(f, \zeta(f)) \succsim (f', \zeta'(f'))$ means that choosing the act f in the first stage followed by implementation of the contingent plan $\zeta(f)$ in the second is at least as preferred as choosing the act f' in the first stage followed by the implementation of the contingent plan $\zeta'(f')$ in the second. The strict preference relation, \succ , and the indifference relation, \sim , are the asymmetric and symmetric parts of \succsim , respectively. A preference relation, \succsim , is *nontrivial* if the corresponding strict preference relation is non-empty. I assume throughout that the preference relations being considered are nontrivial. The two components of \mathbb{C} are *essential* if $\neg((f, \zeta(f)) \sim (f', \zeta'(f'))), \forall \zeta(f), \zeta'(f') \in F^X$ and $f \in F$ and $\neg((f, \bar{\zeta}(f)) \sim (f', \bar{\zeta}'(f'))), \forall f, f' \in F$ such that $\bar{\zeta}(f) = \bar{\zeta}'(f')$.

3. Preference relations: structures and representation

3.1. The axiomatic structure

The first two axioms are standard and require no elaboration.

(A.1) **(Weak Order)** \succsim on \mathbb{C} is complete and transitive.

(A.2) **(Archimedean)** For each $(f, \zeta(f)), (f', \zeta'(f')), (f'', \zeta''(f'')) \in \mathbb{C}$ such that $(f, \zeta(f)) \succ (f', \zeta'(f')) \succ (f'', \zeta''(f''))$ there are $\alpha, \beta \in (0, 1)$ such that $\alpha(f, \zeta(f)) + (1 - \alpha)(f'', \zeta''(f'')) \succ (f', \zeta'(f')) \succ \beta(f, \zeta(f)) + (1 - \beta)(f'', \zeta''(f''))$.

The third axiom is the independence axiom of expected utility theory applied to \mathbb{C} . It asserts the separability inherent in a preference relation that ranks two mixtures of act-contingent-plan pairs, independently of act-contingent-plan pair that is common in the two mixtures.

⁷ Extension to finite sequences complicates the analysis without offering new insights, because, at each stage, the decision maker accumulates a finite history of observations describing the actions chosen and outcomes obtained in the preceding stages, the sequence of posterior distributions on the set of theories evolve as a function of these histories.

(A.3) (**Independence**) For all $(f, \zeta(f)), (f', \zeta'(f')), (f'', \zeta''(f'')) \in \mathbb{C}$ and $\alpha \in (0, 1]$, $(f, \zeta(f)) \succcurlyeq (f', \zeta'(f'))$ if and only if $\alpha (f, \zeta(f)) + (1 - \alpha) (f'', \zeta''(f'')) \succcurlyeq \alpha (f', \zeta'(f')) + (1 - \alpha) (f'', \zeta''(f''))$.

Independence holds for each component separably.⁸

The next axiom formalizes the idea that theories, being abstract ideas, inform the decision-making process by predicting the outcomes of acts but do not affect the decision-maker's well-being directly. To state this idea formally, it is necessary to introduce additional notation and definitions.

Given any two acts, f and f' , let the act $(f_{-t}f')$ be constructed by replacing the t -th coordinate of f by that of f' . Formally, given $f, f' \in F$, define $(f_{-t}f') \in F$ as follows: $(f_{-t}f')(t') = f(t')$ if $t' \in T \setminus \{t\}$ and $(f_{-t}f')(t) = f'(t)$ if $t' = t$. Similarly, given $\zeta(f), \zeta'(f) \in \mathcal{Z}(f)$, let $\zeta(f)_{-t}\zeta'(f) \in \mathcal{Z}(f)$ be the contingent plan obtained by replacing the t -th coordinate of the contingent plan $\zeta(f)$ by that of $\zeta'(f)$. Formally, for every given $\zeta(f), \zeta'(f) \in \mathcal{Z}(f)$ and $x \in X$, $(\zeta(f, x)_{-t}\zeta'(f, x))(t') = \zeta(f, x)(t')$ if $t' \in T \setminus \{t\}$ and $(\zeta(f, x)_{-t}\zeta'(f, x))(t) = \zeta'(f, x)(t)$ if $t' = t$.

A theory is irrelevant if the decision-maker believes that the theory is invalid and, consequently, insofar as the evaluation of the acts and contingent plans is concerned, may be disregarded. Formally, a theory t is *ex ante irrelevant* if, for all $f, f', f'' \in F$, $(f_{-t}f', \zeta(f_{-t}f')) \sim (f_{-t}f'', \zeta(f_{-t}f''))$ for all $\zeta(f_{-t}f'), \zeta(f_{-t}f'')$ such that $\zeta(f_{-t}f', x) = \zeta(f_{-t}f'', x)$, for all $x \in X$. A theory, t , is *ex post irrelevant* if it is *ex ante irrelevant* or if $x \notin X(t, f)$ and is *ex post relevant* if $x \in X(t, f)$.⁹

Two acts, $f, f' \in F$ are said to *agree outside t* if $f_{-t} = f'_{-t}$. Similarly, two contingent plans $\zeta(f), \zeta'(f) \in \mathcal{Z}(f)$ are said to *agree outside t* if $\zeta(f, x)_{-t} = \zeta'(f, x)_{-t}$, for all $x \in X$. The next axiom asserts that if the predictions of any two relevant theories are the same, then the preference ranking of any two alternatives in \mathbb{C} that agree outside one of these theories is the same as that of any two alternatives that agree outside the other theory. This assertion is formulated separately for *ex ante* and *ex post* relevant theories. Formally,

(A.4) (**Theory-independence**) (a) For all $f, f', f'' \in F$, $\zeta, \zeta' \in \mathcal{Z}$ and *ex ante* relevant $t, t' \in T$, if $f'(t) = f'(t')$, $f''(t) = f''(t')$, $\zeta(f_{-t}f') = \zeta'(f_{-t}f'')$ and $\zeta(f_{-t'}f') = \zeta'(f_{-t'}f'')$ then $(f_{-t}f', \zeta(f_{-t}f')) \succcurlyeq (f_{-t}f'', \zeta'(f_{-t}f''))$ if and only if $(f_{-t'}f', \zeta(f_{-t'}f')) \succcurlyeq (f_{-t'}f'', \zeta'(f_{-t'}f''))$. (b) For all $f \in F$, $\zeta, \zeta' \in \mathcal{Z}$, and *ex post* relevant $t, t' \in T$, $(f, \zeta(f, x)_{-t}\zeta'(f, x)) \succcurlyeq (f, \zeta(f, x)_{-t'}\zeta'(f, x))$ if and only if $(f, \zeta(f, x)_{-t'}\zeta'(f, x)) \succcurlyeq (f, \zeta(f, x)_{-t'}\zeta'(f, x))$, for all $x \in X$.

3.2. Representation

The first result characterizes the existence and uniqueness of subjective expected utility representation of \succcurlyeq on \mathbb{C} . The subjectivity is the decision-makers' idiosyncratic valuations of the sequences of outcomes and their evolving personal beliefs regarding the validity of the theo-

⁸ Independence implies that, for all $(f, \zeta(f)), (f, \zeta'(f)), (f, \zeta''(f)) \in \mathbb{C}$ and $\alpha \in (0, 1]$, $(f, \zeta(f)) \succcurlyeq (f, \zeta'(f))$ if and only if $(f, \alpha \zeta(f) + (1 - \alpha) \zeta''(f)) \succcurlyeq (f, \alpha \zeta'(f) + (1 - \alpha) \zeta''(f))$ and for all $(f, \bar{\zeta}(f)), (f', \bar{\zeta}'(f')), (f'', \bar{\zeta}''(f'')) \in \mathbb{C}$ such that $\bar{\zeta}(f) = \bar{\zeta}'(f') = \bar{\zeta}''(f'') = \hat{f}$ and $\alpha \in (0, 1]$, $(f, \hat{f}) \succcurlyeq (f', \hat{f})$ if and only if $(\alpha f + (1 - \alpha) f'', \hat{f}) \succcurlyeq (\alpha f' + (1 - \alpha) f'', \hat{f})$.

⁹ This is analogous to the notion of null events in the theory of decision making under uncertainty.

ries. The joint probability on the sequences of outcomes is induced by the prior and posterior subjective probabilities on the set of theories.

Theorem 1. A preference relation \succsim on \mathbb{C} is nontrivial Archimedean weak order satisfying independence and theory-independence if and only if there is nonconstant function $u : X \times X \rightarrow \mathbb{R}$, and a probability distribution η on T such that, for all $(f, \zeta(f)), (f', \zeta'(f')) \in \mathbb{C}$,

$$(f, \zeta(f)) \succsim (f', \zeta'(f'))$$

if and only if

$$\begin{aligned} & \sum_{x \in X} \left[\sum_{x' \in X} u(x, x') \sum_{t \in T} \mu_t \left(\zeta(f, x)^{-1}(x') \right) \eta(t | f, x) \right] \sum_{t \in T} \mu_t \left(f^{-1}(x) \right) \eta(t) \quad (1) \\ & \geq \sum_{x \in X} \left[\sum_{x' \in X} u(x, x') \sum_{t \in T} \mu_t \left(\zeta'(f', x)^{-1}(x') \right) \eta(t | f', x) \right] \sum_{t \in T} \mu_t \left(f'^{-1}(x) \right) \eta(t). \end{aligned}$$

Moreover, u is unique up to positive affine transformation, and η is unique.

The representation in Theorem 1 may be reformulated as follows: Let $\Pr(x | f)$ and $\Pr(x' | \zeta(f, x))$ denote the prior distribution of X given f and the posterior distribution on X given the observation (f, x) respectively. Formally,

$$\Pr(x | f) := \sum_{t \in T} \mu_t \left(f^{-1}(x) \right) \eta(t) \quad (2)$$

and

$$\Pr(x' | \zeta(f, x)) := \sum_{t \in T} \mu_t \left(\zeta(f, x)^{-1}(x') \right) \eta(t | f, x), \quad \forall (f, x) \in F \times X. \quad (3)$$

Then the joint probability distribution on $X \times X$ induced by $(f, \zeta(f))$ is:

$$\Pr(x, x' | (f, \zeta(f))) = \Pr(x' | \zeta(f, x)) \Pr(x | f). \quad (4)$$

The representation may be expressed as follows:

$$(f, \zeta(f)) \mapsto \sum_{x \in X} \sum_{x' \in X} u(x, x') \Pr(x, x' | f, \zeta(f)). \quad (5)$$

The decision-maker is *Bayesian* if the second-stage conditional probabilities are $\eta(t | f, x) = \mu_t(f^{-1}(x)) \eta(t) / \sum_{t' \in T} \mu_{t'}(f^{-1}(x)) \eta(t')$, for all $(f, x) \in F \times X$ and $t \in T$. In this case, (2) and (4) imply that

$$\Pr(x, x' | f, \zeta(f)) = \sum_{t \in T} \mu_t \left(\zeta(f, x)^{-1}(x') \right) \mu_t \left(f^{-1}(x) \right) \eta(t), \quad \forall (f, x) \in F \times X.$$

In particular, $\eta(t) = 0$ if and only if t is ex-ante irrelevant and $\eta(t | f, x) = 0$ if and only if t is ex poste irrelevant.

A special case of the representation Theorem 1 is the additive representation. To obtain such representation, recall that F and $\mathcal{Z}(f)$ are convex sets in $\mathbb{R}^{|T \times X|}$ and $\mathbb{R}^{|T \times X \times X|}$, respectively. Suppose that they are endowed with the topology of \mathbb{R}^n , and the choice set, \mathbb{C} , is endowed with the product topology.

Assume that the two components of \mathbb{C} are *essential* and that the preference relation \succsim on \mathbb{C} satisfies the following, well-known, hexagon condition.

(A.5) (**Hexagon condition**): For all $f, f', f'' \in F$ and $\zeta, \zeta', \zeta'' \in \mathcal{Z}(f, \zeta'(f)) \sim (f', \zeta(f'))$ and $(f, \zeta''(f)) \sim (f', \zeta'(f')) \sim (f'', \zeta(f''))$ then $(f', \zeta''(f')) \sim (f'', \zeta'(f'))$.

The preference relation is *continuous* if the upper and lower contour sets, $\{(f', \zeta(f')) \in \mathbb{C} \mid (f', \zeta(f')) \succsim (f, \zeta(f))\}$ and $\{(f', \zeta(f')) \in \mathbb{C} \mid (f, \zeta(f)) \succsim (f', \zeta(f'))\}$, are closed for all

$(f, \zeta(f)) \in \mathbb{C}$. A pair of functions, (u_1, u_2) are jointly cardinal representation of \succsim if the class of all such pairs (v_1, v_2) that represent \succsim satisfy: $v_i = bu_i + a_i, b > 0, i = 1, 2$.

With this in mind we can state the following result:

Theorem 2. *A preference relation \succsim on \mathbb{C} is a nontrivial continuous weak order satisfying independence, theory-independence, and the hexagon condition if and only if there are nonconstant, real-valued functions u_1 and u_2 on X and a probability distribution η on T such that, for all $(f, \zeta(f)), (f', \zeta'(f')) \in \mathbb{C}$,*

$$(f, \zeta(f)) \succsim (f', \zeta'(f'))$$

if and only if

$$\begin{aligned} \sum_{x \in X} u_1(x) \Pr(x | f) + \sum_{x' \in X} u_2(x') \sum_{x \in X} \Pr(x' | \zeta(f, x)) \Pr(x | f) &\geq \\ \sum_{x \in X} u_1(x) \Pr(x | f) + \sum_{x' \in X} u_2(x') \sum_{x \in X} \Pr(x' | \zeta'(f', x)) \Pr(x | f). \end{aligned}$$

Moreover, u_1 and u_2 are jointly cardinal and η is unique.

In general, decisions involve trade-offs between material benefits and information acquisition. The model described here allows for an act to be chosen that foregoes imminent material benefits if it generates information that improves the subsequent choices. This point is illustrated by the following example.

Exploitation-exploration trade-off: Consider the urn example described in the introduction. There are two hypotheses regarding the process. According to hypothesis *I*, balls are drawn with replacement and according to hypothesis *II*, balls are drawn without replacement. Let the set of payoffs be $X = \{\$0, \$x, \hat{x}\}$, and consider a choice between two acts f_1 is a bet on the outcome of a draw of two balls from the urn that pays x dollars if the two balls are of the same color and zero dollars otherwise, and f_2 is a bet on the outcome of a draw of three balls from the urn that pays \hat{x} dollars if they are of the same color and zero otherwise. Assume that the payoffs, $x > 0$ and $\hat{x} > 0$, are such that

$$x \Pr(x | f_1) = \hat{x} \Pr(\hat{x} | f_2), \tag{6}$$

where

$$\Pr(x | f_1) = \mu_I(f_1^{-1}(x)) \eta(I) + \mu_{II}(f_1^{-1}(x)) \eta(II)$$

and

$$\Pr(\hat{x} | f_2) = \mu_I(f_2^{-1}(\hat{x})) \eta(I) + \mu_{II}(f_2^{-1}(\hat{x})) \eta(II).$$

Suppose that the utility function is $u(x, x') = x + x'$. Denote by $\zeta^*(f, z)$ the solution to

$$\max_{f \in \{f_1, f_2\}} \sum_{x' \in \{\hat{x}, x, 0\}} x' \Pr(x' | f, z), \quad z \in \{0, x\}.$$

Then the representation (5) implies that choosing f_1 in the first stage and proceeding optimally yields

$$x \Pr(x | f_1) + \sum_{z \in \{x, 0\}} \sum_{x' \in \{\hat{x}, x, 0\}} x' \Pr(x' | \zeta^*(f_1, z)) \Pr(z | f_1),$$

and choosing f_2 in the first stage and proceeding optimally yields

$$\hat{x} \Pr(\hat{x} \mid f_2) + \sum_{z \in \{\hat{x}, 0\}} \sum_{x' \in \{\hat{x}, x, 0\}} x' \Pr(x' \mid \zeta^*(f_2, z)) \Pr(z \mid f_2).$$

By (6), the exploitation value (i.e., the first-stage expected payoff) is the same under the two acts. However, according to the definition of Blackwell (1951), f_2 is *sufficient* for f_1 .¹⁰ Hence, by Blackwell's (1951) theorem, f_2 is more informative, that is:

$$\begin{aligned} & \sum_{z \in \{\hat{x}, 0\}} \sum_{x' \in \{\hat{x}, x, 0\}} x' \Pr(x' \mid \zeta^*(f_2, z)) \Pr(z \mid f_2) \\ & > \sum_{z \in \{x, 0\}} \sum_{x' \in \{\hat{x}, x, 0\}} x' \Pr(x' \mid \zeta^*(f_1, z)) \Pr(z \mid f_1). \end{aligned}$$

Consequently, f_2 has a higher exploration value.

Thus, choosing f_2 in the first stage and proceeding optimally is preferred over choosing f_1 in the first stage and proceeding optimally from there. By continuity of the first stage expected payoff, for $\varepsilon > 0$ sufficiently small, an increase from x to $x + \varepsilon$ makes f_1 strictly better than f_2 from the exploitation viewpoint but does not reverse the preference ranking (that is, $(f_2, \zeta^*(f_2))$ is strictly preferred over $(f_1, \zeta^*(f_1))$).

4. Experiments

4.1. Preliminaries

Experiments are acts the outcomes of which, dubbed *signals*, are devoid of material implications. Formally, experiments are random variables, \tilde{y} , on the Borel-measurable space Ω taking values in a set of signals, $Y \subset X$. Thus, the set \mathcal{E} of experiments is a subset of F . Let $\mathcal{Z} = \{\zeta : \mathcal{E} \times Y \rightarrow F\}$ be sets of mappings representing plans of choosing acts contingent on the observations. Let $\mathcal{Z}(\tilde{y}) := \{\zeta(\tilde{y}) : Y \rightarrow F \mid \zeta \in \mathcal{Z}\}$. Then the choice set is $\mathcal{C} := \{(\tilde{y}, \zeta(\tilde{y})) \in \mathcal{E} \times \mathcal{Z}(\tilde{y})\}$.

A central tenet of the subjective expected utility theory is that information affects decision-makers' beliefs while leaving their tastes intact. To formalize this premise, I propose a variation of the model of the preceding section in which the first-stage decision is a choice of an experiment, \tilde{y} , followed, in the second stage, by a choice of an act contingent on the observations $(\tilde{y}, y) \in \mathcal{E} \times Y$. The idea is that, based on the observation obtained in the first stage, the decision-maker updates his beliefs about the validity of the underlying theories and, consequently, his preferences over the second-stage acts.

To formalize the idea that neither the experiment itself nor its signals affects the decision-maker's well-being except through the update of his beliefs about the likely outcomes of the second-stage acts, it is necessary to separate the informational effects of experiments from potential signal effects on the decision-maker's well-being.

4.2. Preferences and representation

The idea that decision-makers regard experiments as valuable only inasmuch as they are informative, is captured by the requirement that information that is not exploitable is valueless and, consequently, the experiments that generate it belong to the same indifference class. For instance,

¹⁰ To see this, let the information structures corresponding to f_1 and f_2 be given by the right stochastic matrices $Q_1 = \begin{pmatrix} 1/2 & 1/2 \\ 49/99 & 50/99 \end{pmatrix}$ and $Q_2 = \begin{pmatrix} 1/4 & 3/4 \\ 49 \times 48/99 \times 98 & 1 - 49 \times 48/99 \times 98 \end{pmatrix}$. Then $Q_2 M = Q_1$, where the garbling matrix M is: $\begin{pmatrix} 51/99 & 48/99 \\ 49/99 & 50/99 \end{pmatrix}$.

if the feasible set of acts is a singleton (i.e., there is no choice to speak of), then experimentation is useless and all experiments are equally (non)valuable.

The next axiom states this assertion formally. It requires that all experiments followed by contingent plans that do not permit the exploitation of information, are indifferent to one another. Stating the axiom formally requires the following additional notation and definitions: For each $\tilde{y} \in \mathcal{E}$ the support of \tilde{y} is $S(\tilde{y}) = \{y \in Y \mid \sum_{t \in T} \mu_t(\tilde{y}^{-1}(y)) \eta(t) > 0\}$. For every $\tilde{y} \in \mathcal{E}$, a *constant contingent plan* is $\zeta(\tilde{y}) \in \mathcal{Z}(\tilde{y})$ such that $\zeta(\tilde{y}, y) = \zeta(\tilde{y}, y')$, for all $y, y' \in S(\tilde{y})$. The set of the constant contingent plans is identified with F . Invoking this identification, the axiom asserts that $(\tilde{y}, f) \sim (\tilde{y}', f)$, for all $\tilde{y}, \tilde{y}' \in \mathcal{E}$ and $f \in F$.

(A.6) (**Signal-independence**) For all $\tilde{y}, \tilde{y}' \in \mathcal{E}$ and constant contingent plans $\zeta(\tilde{y}) = f = \zeta(\tilde{y}')$, $(\tilde{y}, f) \sim (\tilde{y}', f)$.

The next theorem characterizes the representation of preference ranking of experiments. It is obtained from Theorem 1 by restricting choice space to alternatives that are devoid of material consequences and amending the structure of preference relation with signal-independence. Unlike Theorem 1, according to which the ranking of alternatives in \mathbb{C} is based on their exploitation and exploration values, the ranking of experiments is motivated solely by their exploration values.

Theorem 3. *A binary relation \succsim on \mathcal{C} is a nontrivial, Archimedean, weak order satisfying independence, theory-independence, and signal-independence if and only if there is a non-constant function $u : X \rightarrow \mathbb{R}$, and probability distribution η on T such that for all $(\tilde{y}, \zeta(\tilde{y})), (\tilde{y}', \zeta'(\tilde{y}')) \in \mathcal{C}$,*

$$(\tilde{y}, \zeta(\tilde{y})) \succsim (\tilde{y}', \zeta'(\tilde{y}'))$$

if and only if

$$\begin{aligned} & \sum_{y \in S(\tilde{y})} \left[\sum_{x \in X} u(x) \sum_{t \in T} \mu_t \left(\zeta(\tilde{y}, y)^{-1}(x) \right) \eta(t \mid \tilde{y}, y) \right] \sum_{t \in T} \mu_t \left(\tilde{y}^{-1}(y) \right) \eta(t) \\ & \geq \sum_{y \in S(\tilde{y}')} \left[\sum_{x \in X} u(x) \sum_{t \in T} \mu'_t \left(\zeta'(\tilde{y}', y)^{-1}(x) \right) \eta'(t \mid \tilde{y}', y) \right] \sum_{t \in T} \mu_t \left(\tilde{y}'^{-1}(y) \right) \eta'(t). \end{aligned}$$

Moreover, u is unique up to positive affine transformation, and η is unique.

5. Elicitation of the subjective probabilities

5.1. The elicitation problem

Incentive-compatible mechanisms designed to elicit subjective probabilities on a state space have been studied for more than half a century. Pioneered by the works of Brier (1950) and Good (1952), these studies include Savage (1971), Grether (1981), Kadane and Winkler (1988), and Karni (2009).¹¹ A common feature of these elicitation schemes is the conditioning of the subject's reward on the events of interest. This conditioning requires that the occurrence of the events of interest be verifiable. Because, in general, theories are not verifiable, these mechanisms do not apply to the elicitation of a subject's prior beliefs about the truth of theories.

Prelec (2004); Chambers and Lambert (2015, 2021); and Karni (2020) proposed incentive-compatible mechanisms designed to elicit subjective probabilities of events the occurrence of

¹¹ For a recent comprehensive review, see Chambers and Lambert (2021).

which is private information and, consequently, unverifiable. However, the working of these mechanisms hinges on the presumption that the subject discovers, for himself, the truth of the unobservable event of interest. Because the uncertainty about the truth of theories may not dissipate in the subject’s own mind, these mechanisms do not apply to the elicitation problem with which we are concerned.

Invoking the observability of the signals of experiments, I propose next a new, indirect, incentive-compatible scheme designed to elicit the subjective probabilities representing the subject’s degree of belief in the truth of the theories and examine the conditions under which it yields the desired outcome.

5.2. The elicitation mechanism

Consider an experiment $\tilde{y} \in \mathcal{E}$ whose support, $S(\tilde{y})$, has cardinality that is at least as great as that of the set of theories, T . Let $\Upsilon(\tilde{y}) = (Y_1, \dots, Y_{|T|})$ be a partition of $S(\tilde{y})$ and denote by $\Upsilon_{|T|}$ the set of all partitions of $S(\tilde{y})$ that have $|T|$ cells.¹²

Fix $\Upsilon(\tilde{y}) \in \Upsilon_{|T|}$. Since the signals are verifiable, it is possible to apply one of the existing schemes (e.g., Karni (2009)) to elicit the subject’s subjective probabilities of the cells of the partition, $P(Y_i)$, $i = 1, \dots, |T|$. By Theorem 3, for all $Y_i \in \Upsilon(\tilde{y})$, $P(Y_i) = \sum_{y \in Y_i} \sum_{t \in T} \mu_t(\tilde{y}^{-1}(y)) \eta(t)$.

For each $Y_i \in \Upsilon(\tilde{y})$ and $t \in T$, let $\xi_t(Y_i) := \sum_{y \in Y_i} \mu_t(\tilde{y}^{-1}(y))$ (i.e., $\xi_t(Y_i)$ denotes the probability that the theory t assigns to the set of signals Y_i conditional on the experiment \tilde{y}). Define $\eta(\Upsilon(\tilde{y})) = (\eta(t_1 | \Upsilon(\tilde{y})), \dots, \eta(t_{|T|} | \Upsilon(\tilde{y})))$. Then $A\eta^\tau(\Upsilon(\tilde{y})) = (P(Y_1), \dots, P(Y_{|T|-1}), 1)^\tau$, where the superscript τ is the transpose and A is the $|T| \times |T|$ matrix given by:

$$A = \begin{bmatrix} \xi_{t_1}(Y_1) & \cdot & \cdot & \xi_{t_{|T|}}(Y_1) \\ \cdot & \cdot & \cdot & \cdot \\ \xi_{t_1}(Y_{|T|-1}) & \cdot & \cdot & \xi_{t_{|T|}}(Y_{|T|-1}) \\ 1 & 1 & 1 & 1 \end{bmatrix}. \tag{7}$$

The following proposition is immediate:

Proposition. *The probability distribution $\eta(\Upsilon(\tilde{y}))$ on T exists and is unique if and only if there is an experiment $\tilde{y} \in \mathcal{E}$ and a partition of $S(\tilde{y})$ such that the corresponding matrix A is nonsingular.*

Assume that for every $\Upsilon(\tilde{y}) \in \Upsilon_{|T|}$, $\eta(\Upsilon(\tilde{y}))$ exists. A prior probability distribution η on T is said to *represent a decision-maker’s beliefs* if and only if

$$\eta(\Upsilon(\tilde{y})) = \eta(\Upsilon'(\tilde{y})) = \eta(\Upsilon(\tilde{y}')) = \eta(\Upsilon'(\tilde{y}')) = \eta,$$

for all $\Upsilon(\tilde{y}), \Upsilon'(\tilde{y}), \Upsilon(\tilde{y}'), \Upsilon'(\tilde{y}') \in \Upsilon_{|T|}$ and $\tilde{y}, \tilde{y}' \in \mathcal{E}$.¹³

¹² If $|S(\tilde{y})| = |T|$, then cells of the partition are singleton sets, each containing an element of the support of \tilde{y} .
¹³ If no single experiment has support whose cardinality is at least as great as that of the set of relevant theories but $\sum_{\tilde{y} \in \mathcal{E}} |S(\tilde{y})| \geq |T|$, it is possible to run a number of experiments, partition the sets of signals so that the total number of cells is equal to $|T|$ and elicited the corresponding vectors of probabilities, \hat{P} . A unique subjective probability η on T is given by the solution to the system of equations $\hat{A}\eta^\tau = \hat{P}^\tau$ if and only if the matrix \hat{A} , whose columns are the probabilities assigned by the different theories to the cells of the partition is nonsingular.

It is important to underscore that the proposed elicitation mechanism may not work if it is applied to general acts, because the valuations of the payoffs are not generally outcome-independent. Consequently, the elicitation mechanisms fail to produce reliable values of the probabilities of the cells of the partition of the set of outcomes.¹⁴

6. Discussion

6.1. States and consequences

Perhaps the simplest way to explain the distinction between the traditional analytical framework employed in the theories of decision making under uncertainty and the one proposed here is by reviewing the similarities to and differences with the model of Anscombe and Aumann (1963). Because the theory-dependent consequences of acts are distributions on X , the decision model of this paper may seem analogous to that of Anscombe and Aumann with theories replacing the states. In particular, like the states in the framework of Anscombe and Aumann, theories are exogenously given. However, unlike in the Anscombe-Aumann model, in which states are verifiable, in general, theories are not. In other words, the Anscombe-Aumann states – whether they are the outcomes of actions, as in the case of running a horse race, or occur naturally, as in the case of the weather – are verifiable independently of the acts.¹⁵ By contrast, because observations may be consistent with distinct theories, in general observing an act-outcome pair may be sufficient to invalidate a theory but not sufficient to validate it.¹⁶

Another difference has to do with the interpretation of the subjective probabilities. Unlike the Anscombe-Aumann model – in which decision-makers entertain beliefs, represented by subjective probabilities, on the likelihoods of states (e.g., the outcomes of a horse race or the weather) – in the model of this paper, these states are outcomes, and their probabilities are objective forecasts of the underlying theories (e.g., “horse race theories” or weather forecast models). Instead of beliefs on states, decision-makers’ entertain beliefs about the truth of the alternative theories.

6.2. Multi-period extension

Consider the extension of the model to include a finite set, $T = \{0, 1, \dots, n\}$, of decision nodes, where 0 denotes the first (root) choice node. At each decision node the decision maker chooses an act, $f \in F$, following which he is awarded an outcome $x \in X$, selected randomly. The extensive form of the decision making process may be depicted by a tree with alternating decision nodes (nodes at which acts are chosen) and chance nodes (nodes at which outcomes are realized).

The choice set is: $\mathbb{C} = F \times \mathcal{Z}$, where \mathcal{Z} denotes the set of plans for choosing subsequent acts contingent on the histories of observations (i.e., sequences of act-outcome pairs). A typical element of $\zeta \in \mathcal{Z}$ is constructed as follows: Given $f^0 \in F$, let $f^1 = \zeta(f^0, x^1)$, $f^2 = \zeta(f^0, \zeta(f^0, x^1), x^2)$, ..., $f^j = \zeta(h^j, x^j)$, ..., $f^n = \zeta(h^{n-1}, x^n)$, where $x^j \in X$ and, for every history of observations, $h^j := (f^0, (f^0, x^1), \dots, (f^{j-1}, x^j)) \in (F \times X)^j$, $j = 1, \dots, n - 1$. For example, if $n = 3$, we have:

¹⁴ This is a version of the familiar state-dependent problem with the known elicitation procedures.

¹⁵ This is a necessary condition for the payoffs of the acts to be effectuated.

¹⁶ One may bet on the falsification of a theory by observations. One may not bet, however, on a theory being proved to be true.

$$\zeta = \left(\zeta \left(f^0, x^1 \right), \zeta \left(f^0, \zeta \left(f^0, x^1 \right), x^2 \right) \right), \\ \zeta \left(f^0, \zeta \left(f^0, x^1 \right), \zeta \left(f^0, \zeta \left(f^0, x^1 \right), x^2 \right), x^3 \right).$$

Then, for all $(f^0, \zeta) \in \mathbb{C}$, the representation of Theorem 1 is given by:

$$(f^0, \zeta) \mapsto \sum_{(x^1, x^2, x^3) \in X^3} u \left((x^1, x^2, x^3) \right) \Pr \left((x^1, x^2, x^3) \mid f^0, \zeta \right),$$

where

$$\Pr \left((x^1, x^2, x^3) \mid f^0, \zeta \right) \\ = \Pr \left(x^1 \mid f^0 \right) \Pr \left(x^2 \mid f^0, f^1 \left(x^1 \right) \right) \Pr \left(x^3 \mid f^0, f^1 \left(x^1 \right), f^2 \left(x^2 \right) \right).$$

Applying of Bayes' rule, $\Pr \left((x^1, x^2, x^3) \mid f^0, \zeta \right)$ is:

$$\sum_{t \in T} \mu_t \left(\left(f^0 \right)^{-1} \left(x^1 \right) \right) \mu_t \left(\zeta^{-1} \left(f^0, \zeta \left(f^0, x^1 \right) \right) \left(x^2 \right) \right) \\ \times \mu_t \left(\zeta^{-1} \left(f^0, \zeta \left(f^0, x^1 \right), \zeta \left(f^0, \zeta \left(f^0, x^1 \right), x^2 \right) \right) \left(x^3 \right) \right) \eta(t).$$

Extension of Theorem 2 with coordinate independence replacing the hexagon condition is straightforward.¹⁷

6.3. Related literature

Predicting the outcomes of acts on the basis of laws that capture the regularity of the relation between acts and outcomes seems to conform to the way we think. It also conforms with the scientific method, according to which general laws are parsimonious and efficient means of describing the environment relevant to a decision problem. The notion that choice under uncertainty is based on laws that assign probability distributions on the state space was suggested by Klibanoff et al. (2005) and was adopted by Denti and Pomatto (2021). In both cases, this notion is offered as a possible interpretation of the elements of the set of priors that figure in their respective models of smooth ambiguity. Cerreia-Vioglio et al. (2021) employ a similar idea, suggesting that decision-makers invoke “structured models” to assess the uncertainty induced by model misspecification. However, apart from this feature, the models of smooth ambiguity and model misspecification and the model of this paper are different in their objectives and, consequently, the analytical frameworks they employ and the structures of the preference relations. Perhaps the most important difference is that the aforementioned models do not include the dynamics that are at the core of the exploitation-exploration process modeled in this paper.

Hyogo (2007) proposes a different decision theoretical model of experimentation, the focal point of which is the subjective interpretation of relation between experiments and the distribution of signals. Decisions in Hyogo’s model span two periods. In the first period, the decision-maker is supposed to choose an action and a subset of Anscombe-Aumann acts that is referred to as a menu. The action generates a signal, which the decision-maker uses to update her beliefs about the likelihoods of the states. In the second period, the decision-maker chooses an Anscombe-Aumann act from the predetermined menu. An experiment is a pair (Y, I) , where

¹⁷ See Wakker (1989).

$l : S \times A \rightarrow \Delta Y$ is a function, A is the set of actions, S is the set of states of the world, and ΔY is the set of distributions on a set, Y , of signals. The main objective of Hyogo’s model is “to make the pair (Y, l) , in addition to the prior, subjective” (Hyogo (2007), p. 317).

Hyogo’s approach is fundamentally different from the one proposed in this paper in several important respects. The first is the modeling and definition of experiments. The analogue of states of the world in Hyogo’s model are theories and that of actions are random variables on an abstract measure space taking their values in a signal space. However, unlike in Hyogo’s model, the mapping of theory-experiment pairs to the distribution of signals and the set of signals itself are objectively given, because, by definition, a theory generates predictions of the outcomes of experiments. Consequently, the objective of Hyogo’s analysis has no counterpart in the present study, which focuses on the subjective degrees of belief of the decision-maker in the truth of the theories. The different objectives require distinct analytical frameworks. Thus, in Hyogo’s model, elements of choice set in the first period are pairs, consisting of an action and a menu of acts, and that of the second period are acts from the menu that was selected in the first period. In the present model, the elements of the choice set consist of experiments and plans of choosing acts contingent on the experiment-generated signals. Finally, the preference structures and their representations of the two models are different, reflecting the distinct objectives and analytical frameworks.

7. Proofs

7.1. Proof of Theorem 1

(Sufficiency) By the von Neumann-Morgenstern theorem, \succsim is an Archimedean weak order satisfying independence if and only if there is an affine, real-valued, function U on \mathbb{C} such that, for all $(f, \zeta(f)), (f', \zeta'(f')) \in \mathbb{C}$,

$$(f, \zeta(f)) \succsim (f', \zeta'(f')) \Leftrightarrow U(f, \zeta(f)) \geq U(f', \zeta'(f')).$$

Fix $f^* \in F$ and, for any $f \in F$ and $\zeta, \zeta^* \in \mathcal{Z}$. By definition,

$$\frac{1}{m}(f, \zeta(f)) + \frac{m-1}{m}(f^*, \zeta^*(f^*)) = \frac{1}{m} \sum_{t \in T} (f^*_t f(t), \zeta^*(f^*)_{-t} \zeta(f)(t)),$$

where $\zeta(f)(t) = (\zeta(f, x_1)(t), \dots, \zeta(f, x_n)(t)) \in \Delta(X)^n$. By the affinity of U ,

$$\frac{1}{m}U(f, \zeta(f)) + \frac{m-1}{m}U(f^*, \zeta^*(f^*)) = \frac{1}{m} \sum_{t \in T} U(f^*_t f(t), \zeta^*(f^*)_{-t} \zeta(f)(t)).$$

Define a function $W : T \times \Delta X \times (\Delta X)^n \rightarrow \mathbb{R}$ by:

$$W(t, f(t), \zeta(f)(t)) = U\left((f^*_t f(t)), (\zeta^*(f^*)_{-t} \zeta(f)(t))\right) - \frac{m-1}{m}U(f^*, \zeta^*(f^*)).$$

Thus,

$$\begin{aligned} \frac{1}{m} \sum_{t \in T} W(t, f(t), \zeta(f)(t)) &= \frac{1}{m} \sum_{t \in T} U\left((f^*_t f(t)), (\zeta^*(f^*)_{-t} \zeta(f)(t))\right) \\ &\quad - \frac{m-1}{m}U(f^*(t), \zeta^*(f^*)(t)). \end{aligned}$$

Hence,

$$U(f, \zeta(f)) = \sum_{t \in T} W(t, f(t), \zeta(f)(t)).$$

By the affinity of U , $W(t, \cdot, \cdot)$ is affine.

By theory-independence, for all ex-ante relevant $t, t' \in T$ and $f, f' \in F$

$$\left(f_{-t}^* f(t), \zeta^*(f^*)_{-t} \zeta(f)(t)\right) \succcurlyeq \left(f_{-t}^* f'(t), \zeta^*(f^*)_{-t} \zeta(f')(t)\right)$$

if and only if

$$\left(f_{-t'}^* f(t), \zeta^*(f^*)_{-t'} \zeta(f)(t)\right) \succcurlyeq \left(f_{-t'}^* f'(t), \zeta^*(f^*)_{-t'} \zeta(f')(t)\right)$$

Then, by the additivity of U across theories, for all ex-ante relevant $t, t' \in T$,

$$W(t, f(t), \zeta(f)(t)) \geq W(t, f'(t), \zeta(f')(t))$$

if and only if

$$W(t', f(t), \zeta(f)(t)) \geq W(t', f'(t), \zeta(f')(t)).$$

Define $w(\cdot, \cdot) : \Delta X \times (\Delta X)^m \rightarrow \mathbb{R}$ by $w(\cdot, \cdot) = W(t_1, \cdot, \cdot)$.

Then, the uniqueness of the affine utility representation,

$$W(t, \cdot, \cdot) = b_t w(\cdot, \cdot) + a_t,$$

where $b_t > 0$, for all ex-ante relevant $t \in T$ and $b_t = 0$ for all ex-ante irrelevant $t \in T$. By non-triviality, $\sum_{t' \in T} b_{t'} > 0$. Define $\eta(t) = b_t / \sum_{t' \in T} b_{t'}$, for all $t \in T$, then

$$U(f, \zeta(f)) = (\sum_{t' \in T} b_{t'}) \sum_{t \in T} \eta(t) w(f(t), \zeta(f)(t)) + \sum_{t \in T} a_t.$$

Thus,

$$U(f, \zeta(f)) \geq U(f', \zeta'(f'))$$

if and only if

$$\sum_{t \in T} \eta(t) w(f(t), \zeta(f)(t)) \geq \sum_{t \in T} \eta(t) w(f'(t), \zeta'(f')(t)). \tag{8}$$

Let $\delta_x \in F$ assign the outcome x the unit probability mass. Then $\delta_x^{-1} = \Omega$ and $\zeta(\delta_x) \in F$. Define a function $\hat{U} : X \times F \rightarrow \mathbb{R}$ by $\hat{U}(x, \zeta(\delta_x)) = w(\delta_x, \zeta(\delta_x))$, for all $\zeta \in \mathcal{Z}$. By the affinity of w and induction on the size the support (see Kreps (1988) p. 50) we have

$$w(f(t), \zeta(f)(t)) = \sum_{x \in X} \hat{U}(x, \zeta(f, x)(t)) \mu_t(f^{-1}(x)).$$

Hence, by (8),

$$U(f, \zeta(f)) \geq U(f', \zeta'(f'))$$

if and only if

$$\begin{aligned} & \sum_{x \in X} \hat{U}(x, \zeta(f, x)(t)) \sum_{t \in T} \mu_t(f^{-1}(x)) \eta(t) \\ & \geq \sum_{x \in X} \hat{U}(x, \zeta'(f', x)(t)) \sum_{t \in T} \mu_t(f'^{-1}(x)) \eta(t). \end{aligned} \tag{9}$$

Consider next the function $\hat{U}(x, \zeta(f, x)(t))$. Fix $\zeta^* \in \mathcal{Z}$. Then, by the same argument as above,

$$\frac{1}{m} (f, \zeta (f)) + \frac{m-1}{m} (f, \zeta^* (f)) = \frac{1}{m} \sum_{t \in T} (f, \zeta^* (f)_{-t} \zeta (f) (t)).$$

Thus, by (9) and the affinity of $\hat{U} (x, \cdot)$,

$$\begin{aligned} & \sum_{x \in X} \left[\frac{1}{m} \hat{U} (x, \zeta (f, x)) + \frac{m-1}{m} \hat{U} (x, \zeta^* (f, x)) \right] \sum_{t \in T} \mu_t (f^{-1} (x)) \eta (t) \\ &= \sum_{x \in X} \left[\frac{1}{m} \sum_{t' \in T} \hat{U} (x, \zeta^* (f, x)_{-t'} \zeta (f, x) (t')) \right] \sum_{t \in T} \mu_t (f^{-1} (x)) \eta (t). \end{aligned}$$

Define a function $H : T \times X \times \Delta X \rightarrow \mathbb{R}$ by:

$$H (t, x, \zeta (f, x)) = \hat{U} (x, \zeta^* (f, x)_{-t} \zeta (f, x) (t)) - \frac{m-1}{m} \hat{U} (x, \zeta^* (f, x)).$$

Then, by the same argument as above,

$$\hat{U} (x, \zeta (f, x)) = \sum_{t \in T} H (t, x, \zeta (f, x) (t)). \tag{10}$$

By theory independence, for all ex-post relevant $t, t' \in T$ and $\zeta, \zeta' \in \mathcal{Z}$, and for every given $(f, x) \in F \times X$,

$$(f, \zeta^* (f, x)_{-t} \zeta (f, x)) \succcurlyeq (f, \zeta^* (f, x)_{-t} \zeta' (f, x))$$

if and only if,

$$(f, \zeta^* (f, x)_{-t'} \zeta (f, x)) \succcurlyeq (f, \zeta^* (f, x)_{-t'} \zeta' (f, x)).$$

Hence, by (10),

$$\begin{aligned} & \sum_{x \in X} H (t, x, \zeta (f, x) (t)) \sum_{t \in T} \mu_t (f^{-1} (x)) \eta (t) \\ & \geq \sum_{x \in X} H (t, x, \zeta' (f, x) (t)) \sum_{t \in T} \mu_t (f^{-1} (x)) \eta (t) \end{aligned}$$

if and only if

$$\begin{aligned} & \sum_{x \in X} H (t', x, \zeta (f, x) (t)) \sum_{t \in T} \mu_t (f^{-1} (x)) \eta (t) \\ & \geq \sum_{x \in X} H (t', x, \zeta' (f, x) (t)) \sum_{t \in T} \mu_t (f^{-1} (x)) \eta (t). \end{aligned}$$

Define $h (x, \cdot) = H (t_1, x, \cdot)$, then, by the uniqueness of the affine representation, for all $t \in T$,

$$H (t, x, \zeta (f, x) (t)) = \hat{b}_t (f, x) h (x, \zeta (f, x) (t)) + \hat{a}_t (f, x)$$

By nontriviality, $\sum_{t \in T} \hat{b}_t (f, x) > 0$. Let $\eta (t | f, x) := \hat{b}_t (f, x) / \sum_{t' \in T} \hat{b}_{t'} (f, x)$.

Let $\zeta_I (\delta_x) = \delta_x$ for all $x \in X$ and define $u : X \times X \rightarrow \mathbb{R}$ by $u (x, x') = h (x, \zeta_I (\delta_{x'}))$. Then, by the affinity of $h (x, \cdot)$, we have $h (x, \zeta (f, x) (t)) = \sum_{x' \in X} u (x, x') \mu_t (\zeta (f, x)^{-1} (x'))$ and, by (10),

$$\begin{aligned} \hat{U} (x, \zeta (f, x)) &= \sum_{t' \in T} \hat{b}_{t'} (f, x) \sum_{x' \in X} u (x, x') \sum_{t \in T} \mu_t (\zeta (f, x)^{-1} (x')) \eta (t | f, x) \\ & \quad + \sum_{t \in T} \mu_t \hat{a}_t (f, x) \end{aligned} \tag{11}$$

Combining (9) and (11) yields:

$$U (f, \zeta (f)) \geq U (f', \zeta' (f'))$$

if and only if

$$\begin{aligned} & \sum_{x \in X} \left[\sum_{x' \in X} u(x, x') \sum_{t \in T} \mu_t \left(\zeta(f, x)^{-1}(x') \right) \eta(t | f, x) \right] \sum_{t \in T} \mu_t \left(f^{-1}(x) \right) \eta(t) \\ & \geq \sum_{x \in X} \left[\sum_{x' \in X} u(x, x') \sum_{t \in T} \mu_t \left(\zeta'(f', x)^{-1}(x') \right) \eta(t | f, x) \right] \sum_{t \in T} \mu_t \left(f'^{-1}(x) \right) \eta(t). \end{aligned}$$

(Necessity) The necessity of weak order, Archimedean, and independence follow from the von Neumann-Morgenstern theorem. The necessity of theory independence is immediate.

The uniqueness part follows from the uniqueness of U . \square

7.2. Proof of Theorem 2

The necessity is immediate, so I prove the sufficiency part.

Recall that the sets F and $\mathcal{Z}(f)$, $f \in F$ are connected separable topological spaces, the choice set \mathbb{C} is endowed with the product topology, and both components of the elements of \mathbb{C} are essential. Hence, by Wakker (1989) Theorem III.4.1, \succsim is continuous weak order on \mathbb{C} satisfying the hexagon condition then there exist jointly cardinal, continuous additive representation

$$(f, \zeta(f)) \mapsto V_1(f) + V_2(\zeta(f))$$

By independence and theory independence, we have:

$$V_1(f) = \sum_{x \in X} u_1(x) \sum_{t \in T} \mu_t \left(f^{-1}(x) \right) \eta(t),$$

and

$$\begin{aligned} V_2(\zeta(f)) &= \sum_{x' \in X} u_2(x') \sum_{x \in X} \sum_{t \in T} \mu_t \left(\zeta(f, x)^{-1}(x') \right) \eta(t | f, x) \\ &\quad \times \sum_{t \in T} \mu_t \left(f^{-1}(x) \right) \eta(t). \end{aligned}$$

By (2) and (3) $\sum_{t \in T} \mu_t \left(f^{-1}(x) \right) \eta(t) = \Pr(x | f)$ and $\Pr(x' | \zeta(f, x)) := \sum_{t \in T} \mu_t \left(\zeta(f, x)^{-1}(x') \right) \eta(t | f, x)$. Hence,

$$V_1(f) = \sum_{x \in X} u_1(x) \Pr(x | f)$$

and

$$V_2(\zeta(f)) = \sum_{x' \in X} u_2(x') \sum_{x \in X} \Pr(x' | \zeta(f, x)) \Pr(x | f).$$

The joint cardinality of u_1 and u_2 is an implication of the joint cardinality of the additive representation. \square

7.3. Proof of Theorem 3

Proof. By Theorem 1, a preference relation \succsim on \mathcal{C} is nontrivial, continuous weak order satisfying independence and theory-independence if and only if it admits the representation: $(\tilde{y}, \zeta(\tilde{y})) \mapsto U(\tilde{y}, \zeta(\tilde{y}))$, where

$$\begin{aligned} U(\tilde{y}, \zeta(\tilde{y})) &= \sum_{y \in Y} \left[\sum_{x \in X} u(y, x) \sum_{t \in T} \mu_t \left(\zeta(\tilde{y}, y)^{-1}(x) \right) \eta(t | \tilde{y}, y) \right] \\ &\quad \times \sum_{t \in T} \mu_t \left(\tilde{y}^{-1}(y) \right) \eta(t). \end{aligned}$$

I show next that signal-independence is necessary and sufficient condition for u to be independent of the signal, y .

Suppose that $u(y, x) = u(y', x) = u(x)$ for all $y, y' \in Y$, then, for every $\tilde{y} \in \mathcal{E}$ and constant contingent plans $\zeta(\tilde{y}) = f$, it holds that

$$U(\tilde{y}, f) = \sum_{x \in X} u(x) \sum_{t \in T} \mu_t(f^{-1}(x)) \sum_{y \in Y} \eta(t | \tilde{y}, y) \sum_{t \in T} \mu_t(\tilde{y}^{-1}(y)) \eta(t). \quad (12)$$

But $\sum_{y \in Y} \eta(t | \tilde{y}, y) \sum_{t \in T} \mu_t(\tilde{y}^{-1}(y)) \eta(t) = \sum_{y \in Y} \eta(t | \tilde{y}, y) \Pr(y | \tilde{y}) = \eta(t)$. Hence,

$$U(\tilde{y}, f) = \sum_{x \in X} u(x) \sum_{t \in T} \mu_t(f^{-1}(x)) \eta(t).$$

The last expression is independent of \tilde{y} . Thus, $U(\tilde{y}, f) = U(\tilde{y}', f)$, for all $\tilde{y}, \tilde{y}' \in \mathcal{E}$ and $f \in F$. Therefore, by the representation, $(\tilde{y}, f) \sim (\tilde{y}', f)$, for all $\tilde{y}, \tilde{y}' \in \mathcal{E}$ and $f \in F$. Hence, signal independence holds.

Suppose that, for some $y, y' \in Y$, $u(y, \cdot) \neq u(y', \cdot)$. Fix $f \in F$, then $\sum_{t \in T} \mu_t(f^{-1}(x)) \sum_{y \in Y} \eta(t | \tilde{y}, y) = \Pr(x | f)$, for all $\tilde{y}, \tilde{y}' \in \mathcal{E}$. Moreover, for all $\tilde{y} \in \mathcal{E}$, $\sum_{t \in T} \mu_t(\tilde{y}^{-1}(y)) \eta(t) = \Pr(y | \tilde{y})$. Let $\bar{u}(y | f) = \sum_{x \in X} u(x) \Pr(x | f)$ then, by (12),

$$U(\tilde{y}, f) - U(\tilde{y}', f) = \sum_{y \in Y} \bar{u}(y | f) [\Pr(y | \tilde{y}) - \Pr(y | \tilde{y}')].$$

But, by the supposition, $\bar{u}(y | f)$ is not a constant function. Hence, there are $\tilde{y}, \tilde{y}' \in \mathcal{E}$ such that $U(\tilde{y}, f) - U(\tilde{y}', f) \neq 0$. By the representation, $\neg((\tilde{y}, f) \sim (\tilde{y}', f))$, which contradicts signal-independence. Thus, $u(y, \cdot) \neq u(y', \cdot)$ for all $y, y' \in Y$.

The uniqueness follows from Theorem 1. \square

References

Anscombe, Francis J., Aumann, Robert J., 1963. A definition of subjective probability. *Ann. Math. Stat.* 43, 199–205.

Atan, Ontur, Tekin, Cem, van der Schaar, Mihaela, 2015. Global multi-armed bandits with Holder continuity. *Artif. Intell. Stat.* 38, 28–36.

Bergemann, Dirk, Välimäki, Juuso, 2008. Bandit problems. In: Durlauf, Steven N., Blume, Lawrence E. (Eds.), 2nd ed. *The New Palgrave Dictionary of Economics*, vol. 1. Macmillan Press, pp. 336–340.

Blackwell, David, 1951. The comparison of experiments. In: Neyman, J. (Ed.), *Proceedings of the Second Berkley Symposium on Mathematical Statistics and Probability*. University of California, Berkely, pp. 93–102.

Borel, Émile, 1924. A Propos d'un Traité de Probabilités. *Rev. Philos.* 98, 321–336. Translated into English as "Apropos of a Treatise on Probability." Reprinted in *Studies in Subjective Probability*, Henry E. Kyburg Jr. & Howard E. Smokler (Eds.) (1964), Wiley, New York.

Brier, Glenn W., 1950. Verification of forecasts expressed in terms of probabilities. *Mon. Weather Rev.* 78, 1–3.

Cerreia-Vioglio, Simone, Hansen, Lars P., Maccherioni, Fabio, Marinacci, Massimo, 2021. Making Decion under Model Misspecification. Unpublished manuscript.

Chambers, Christopher P., Lambert, Nicolas S., 2015. Introduction to dynamic belief elicitation. *ACM SIGecom Exch.* 14, 80–83.

Chambers, Christopher P., Lambert, Nicolas S., 2021. Dynamic belief elicitation. *Econometrica* 89, 375–414.

Denti, Tommaso, Pomatto, Luciano, 2021. Model and predictive uncertainty: a foundation of smooth ambiguity preferences. *Econometrica*. <https://www.econometricsociety.org>.

Good, Irving J., 1952. Rational decisions. *J. R. Stat. Soc., Ser. B, Methodol.* 14, 107–114.

Grether, David M., 1981. Financial Incentive Effects and Individual Decision Making. *Cal Tech, Social Science Working Paper* 401.

Hyogo, Kazuya, 2007. A subjective model of experimentation. *J. Econ. Theory* 133, 316–330.

Kadane, J.B., Winkler, R.L., 1988. Separating probability elicitation from utility. *J. Am. Stat. Assoc.* 83, 357–363.

Karni, Edi, 2009. A mechanism for eliciting probabilities. *Econometrica* 77, 603–606.

Karni, Edi, 2020. A mechanism for the elicitation of second-order beliefs and subjective information structures. *Econ. Theory* 69, 217–232.

- Klibanoff, Peter, Marinacci, Massimo, Mukerji, Sujoy, 2005. A smooth model of decision making under ambiguity. *Econometrica* 73, 1849–1892.
- Kreps, David, 1988. *Notes on the Theory of Choice*. Westview Press, Boulder and London.
- Pandey, Sandeep, Chakrabarti, Deepayan, Agarwal, Deepak, 2007. Multi-armed bandit problems with dependent arms. In: *Proceedings of the 24th International Conference on Machine Learning*, pp. 721–728.
- Popper, Karl, 2002. *The Poverty of Historicism*. Routledge, London and New York.
- Prelec, Drazen, 2004. A Bayesian truth serum for subjective data. *Science* 306, 462–466.
- Ramsey, Frank P., 1926. *Truth and Probability*. Published in Braithwaite, R.B., Plumpton, F., *The Foundation of Mathematics and Other Logical Essays*, 1931.
- Savage, Leonard J., 1954. *The Foundations of Statistics*. John Wiley, New York.
- Savage, Leonard J., 1971. Elicitation of personal probabilities and expectations. *J. Am. Stat. Assoc.* 66, 783–801.
- Wakker, Peter, 1989. *Additive Representations of Preferences*. Kluwer Academic Publishers, Dordrecht.