A Theory of Subjective Learning[*]

David Dillenberger[†]   Juan Sebastián Lleras[‡]

Philipp Sadowski[§]   Norio Takeoka[¶]

September 2013

Abstract

We study an individual who faces a dynamic decision problem in which the process of information arrival is unobserved by the analyst. We elicit subjective information directly from choice behavior by deriving two utility representations of preferences over menus of acts. One representation uniquely identifies information as a probability measure over posteriors and the other identifies information as a partition of the state space. We compare individuals who expect to learn differently in terms of their preference for flexibility. On the extended domain of dated-menus, we show how to accommodate gradual learning over time by means of a subjective filtration.

Key words: Subjective learning, partitional learning, preference for flexibility, resolution of uncertainty, valuing more binary bets, subjective filtration, dated-menus.

[†]Department of Economics, University of Pennsylvania. E-mail: ddill@sas.upenn.edu

[‡]E-mail: j.sebastian.lleras@gmail.com

[§]Department of Economics, Duke University. E-mail: p.sadowski@duke.edu

[¶]Department of Economics, Yokohama National University. E-mail: takeoka@ynu.ac.jp

1. Introduction

1.1. Motivation and overview

The study of dynamic models of decision making under uncertainty when a flow of information on future risks is expected over time is central in all fields of economics. For example, investors decide when to invest and how much to invest based on what they expect to learn about the distribution of future cash flows. The concepts of value of information and value of flexibility (option value) quantify the positive effects of relying on more precise information structures.[1]

A standard dynamic decision problem has three components: the first component is a set of states of the world that capture all relevant aspects of the decision environment. The second component is a set of feasible intermediate actions, each of which determines the payoff for any realized state. The third component is a description of what the decision maker expects to learn; this component is formalized as an information structure, which is the set of possible signals about the states that are expected to arrive over time and the joint distribution of signals and states.

In many situations, the analyst may be confident in his understanding of the relevant state space and the relevant set of actions. He may, however, not be aware of the information structure people perceive. People may have access to private data which is unforeseen by others; they may interpret data in an idiosyncratic way; or they may be selective in the data they observe, for example by focusing their attention on specific signals. We collectively refer to those situations as "subjective learning". Since information plays a key role in most decision-making processes, if it is not observed, then it should be derived.

In this paper we show how one can infer an individual's subjectively perceived information structure (in addition to his subjective probabilities over states and cardinal utilities over outcomes) solely from his observed choice behavior prior to the resolution of uncertainty. We confine our attention to the study of two canonical ways of describing information, namely identifying signals with the posterior beliefs they produce, and partitional-learning. For each model, we provide an axiomatic foundation and show that the relevant parameters are uniquely identified. The identification of anticipated arrival of information allows us to compare the behavior of individuals who perceive different information structures. Providing a tight link between heterogeneity in information structures and heterogeneity in behavior can guide an applied analyst whose aim is to capture a particular heterogeneity in behavior. Furthermore, such comparisons are the subjective versions of the comparative statics for incremental increases in informativeness when learning is objective. Since the information

---

[1]For a comprehensive survey of the theoretical literature, see Gollier (2001, chapters 24 and 25).

structure is identified from behavior *prior* to the resolution of uncertainty, it may also help an analyst to make a probabilistic forecast of the decision maker's behavior *after* all uncertainty is resolved. Lastly, we propose a new domain that allows studying the behavioral implications of uncertainty that is anticipated to gradually resolve over time by means of a subjective filtration.

The standard subjective expected utility models of Savage (1954) and Anscombe and Aumann (1963) characterize subjective probabilities from observable choices among acts (state-contingent payoffs). These models, however, are not rich enough to also identify un-observed information. The reason is that information has instrumental value only when some aspect of choice can be conditioned on it, which is not the case in the static domain of acts. For example, in environments with potentially asymmetric information, the standard model is not equipped to distinguish which side is better informed. Simultaneously identifying beliefs and information thus requires us to enrich the choice domain, as we now describe.[2]

We consider an objective state space. Actions correspond to acts, that is, state-contingent payoffs, and preferences are defined over sets (or menus) of acts. The interpretation is that the decision maker (henceforth DM) initially chooses among menus and subsequently chooses an act from the menu. If the ultimate choice of an act takes place in the future, then the DM may expect information to arrive prior to this choice. Analyzing today's preferences over future choice situations (menus of acts rather than the acts themselves) allows us to capture the effect of the information the DM expects to learn via his value for flexibility (having more future options available). The preference relation over menus of acts is thus the only primitive of the model, leaving the information structure that the DM faces, as well as his ultimate choice of an act, unmodeled.[3]

To see how our framework can address the question of subjective information acquisition, suppose there are only two states of the world, s_1 and s_2, and consider the acts f, g, and h, given by

$$
\begin{array}{c|cc}
 & s_1 & s_2 \\
f & 1 & 0 \\
g & 0 & 1 \\
h & \frac{3}{4} & \frac{3}{4}
\end{array}
$$

Suppose the DM prefers committing to the act with perfect insurance over either of the bets, that is, $\{h\} \succ \{f\} \sim \{g\}$. Now consider the menu of acts $\{f, g, h\}$. If the DM thinks

---

he might receive a signal that reveals (or is highly indicative of) the true state before he has to choose an act from the menu, then flexibility is valuable and $\{f, g, h\} \succ \{h\}$. If, instead, the DM knows for sure he will not receive such a signal, then $\{f, g, h\} \sim \{h\}$. Similarly, if the agent thinks he might not receive such a signal then $\{f, g, h\} \succ \{f, g\}$, while if he knows for sure that he will receive the signal, then $\{f, g, h\} \sim \{f, g\}$. Thus, what the DM expects to know at the time of choice from the menu is reflected in his preferences over menus of acts.

Section 2 outlines the most general model that captures subjective learning. Theorem 1 derives a *subjective-learning representation* that can be interpreted as follows: the DM behaves as if he has beliefs over the possible posterior distributions over the state space that he might face at the time of choosing from the menu. For each posterior, he expects to choose from the menu the act that maximizes the corresponding expected utility. The model is parameterized by a probability measure on the collection of all possible posterior distributions. This probability measure describes the DM's subjective information structure and is uniquely identified from choice behavior. The axioms that are equivalent to the existence of a subjective-learning representation are the familiar axioms of *Preference relation, vNM Continuity, Nontriviality,* and an appropriate version of *Independence,* in addition to *Domination*, which implies monotonicity in payoffs, and *Set monotonicity*, which captures preference for flexibility.

Identification enables us to compare different decision makers in terms of their preferences for flexibility. We say that DM1 has more preference for flexibility than DM2 if whenever DM1 prefers to commit to a particular action rather than to maintain multiple options, so does DM2. Theorem 2 states that DM1 has more preference for flexibility than DM2 if and only if DM1's distribution of posterior beliefs is a mean-preserving spread of DM2's. This result is analogous to Blackwell's (1951, 1953) comparisons of experiments (in terms of their information content) in a domain where probabilities are objective and comparisons are made with respect to the accuracy of information structures. To rephrase our result in the language of Blackwell, DM1 has more preference for flexibility than DM2 if and only if DM2 would be weakly better off if he could rely on the information structure induced by the subjective beliefs of DM1.

A subjective-learning representation does not allow the identification of information independently of the induced changes in beliefs. The reason is that signals do not correspond to events in the state space. Section 3 addresses this issue by studying the behavioral implications of a (subjective) partitional information structure on a given state space. A partition of the state space is a canonical formalization of information that ensures that the state space

4

captures all the DM expects to know about the decision problem under consideration.[4] This formalization is empirically meaningful: an outside observer who knows the state of the world will also know the information that the decision maker will receive. Theorem 3 derives a *partitional-learning representation* that can be interpreted as follows: the DM has in mind a partition of the state space and prior beliefs over the individual states. The partition describes what he expects to learn before facing the choice of an alternative from the menu. The DM's posterior beliefs conditional on learning an event in the partition are fully determined from the prior beliefs using Bayes' law. For each event, the DM plans to choose an act that maximizes the corresponding expected utility. The partition and the beliefs are endogenous components of the model, which are uniquely identified from choice behavior. Given the assumptions of Theorem 1, Theorem 3 requires only one additional axiom, *Contingent planning*. Suppose the DM knows what information he will receive contingent on the true state of the world. In this case, he is sure about the act he will choose from the menu, contingent on the true state. The DM should then be indifferent between choosing from the menu after learning the deterministic signal, and committing, for every state, to receive the payoff his certain choice would have generated in that state. Axiom Contingent planning captures this indifference.

Individuals who disagree on their prior beliefs are not comparable in terms of their preference for flexibility. Since the description of information as a partition of the state space does not depend on a specific probability distribution, a partitional-learning representation facilitates the behavioral comparisons of such individuals. In particular, partitions can be partially ranked in terms of their fineness, a property which provides a natural measure for the degree of information, independently of any prior beliefs. The behavior of two individuals who expect to receive different information differs in the value they derive from the availability of binary bets. Suppose the DM prefers committing to a constant act that pays c regardless of the state over committing to an act that offers a binary bet on state s versus state s' (in the sense that it pays well on s, pays badly on s', and pays c otherwise). It may still be the case that the DM values the binary bet as an option that is available in addition to c. We say that DM1 *values more binary bets* than DM2 if for any two states s and s' and payoff c for which the premise above holds, whenever DM1 does not value the binary bet as an option in addition to c, neither does DM2. Theorem 4 states that DM1 values more binary bets than DM2 if and only if he expects to receive more information than DM2, in the sense that his partition is finer.[5]

---

[4]For example, it is standard to model information as a partition of the state space in the literature on games with incomplete information, that originated in the seminal contributions of Harsanyi (1967) and Aumann (1976).

[5]An additional requirement for this comparison is that the two DMs agree on whether or not payoffs in

We conclude by studying an individual who anticipates gradual resolution of uncertainty over time. In many applications, the process of information acquisition over time may affect the optimal timing of decisions. For example, an investor may decide to bear the cost of delaying an investment with positive net present value if more information is expected to be received about the distribution of future cash flows. To model such behavior, we introduce the choice domain of *dated-menus* of the form (F, t), where F is a menu and t is the time by which an alternative from the menu must be chosen. Suppose that fixing t, DM's preferences satisfy all the postulates underlying the partitional-learning representation. Furthermore, suppose that the DM prefers to delay his choice from any menu F, in the sense that (F, t) is weakly preferred to (F, t') if $t > t'$, and he is indifferent to the timing if F is a singleton (since then the future choice is trivial). Under these assumptions, the DM has more preference for flexibility at time t than at t', which means, by Theorem 4, that the partition at time t must be finer than that at t'. Using this observation, we provide a *learning by filtration representation*, which suggests that the DM behaves as if he has in mind a filtration indexed by continuous time. Both the filtration, which is the timing of information arrival with the sequence of partitions it induces, and the DM's prior beliefs are uniquely determined from choice behavior.

1.2. Related literature

To our knowledge, very few papers have explored the idea of subjective learning. Dillenberger and Sadowski (2012a) use the same domain as in the present paper to study the most general representation for which signals correspond to events and the DM is Bayesian. They characterize the class of information structures that admit such a representation as a generalization of a set partition, which does not require deterministic signals (that is, the true state of the world may appear in more than one event). Dillenberger and Sadowski show that the generalized-partition model can be applied to study an individual who anticipates gradual resolution of uncertainty over time, without extending the domain as we do in Section 4. Takeoka (2007) uses a different approach to study subjective temporal resolution of uncertainty. He analyzes choice between what one might term "compound menus" (menus over menus etc.) Hyogo (2007) derives a representation that features beliefs over posteriors on a richer domain, where the DM simultaneously chooses a menu of acts and takes an action that might influence the (subjective) process of information arrival.

More generally, our work is part of the preferences over menus literature initiated by

---

state s are valuable at all. This implies that while their prior beliefs need not be the same, they should have the same *support*. If, in addition, their prior beliefs are the same, then DM1 has more preference for flexibility than DM2.

Kreps (1979). Most papers in this literature study uncertainty over future tastes, and not over beliefs on an objective state space. Kreps (1979) studies preferences over menus of deterministic alternatives. Dekel, Lipman, and Rustichini (2001) extend Kreps' domain of choice to menus of lotteries. Our proof of Theorem 1 relies on a sequence of geometric arguments that establish the connection between our domain and theirs. In the setting of preferences over menus of lotteries, Ergin and Sarver (2010) provide an alternative to Hyogo's (2007) approach of modeling costly information acquisition. Recently, De Olivera (2012) and Mihm and Ozbek (2012) combine the framework of the present paper with Ergin and Sarver's idea of costly contemplation to give behavioral foundations to rational inattention.

2. A general model of subjective learning

Let $S = \{s_1, ..., s_k\}$ be a finite state space. An act is a mapping $f : S \rightarrow [0,1]$. Let $\mathcal{F}$ be the set of all acts. Let $\mathcal{K}(\mathcal{F})$ be the set of all non-empty compact subsets of $\mathcal{F}$, endowed with the Hausdorff metric. Capital letters denote sets, or menus, and small letters denote acts. For example, a typical menu is $F = \{f, g, h, ...\} \in \mathcal{K}(\mathcal{F})$. Let $\succeq$ be a binary relation over $\mathcal{K}(\mathcal{F})$. The symmetric and asymmetric components of $\succeq$ are denoted by $\sim$ and $\succ$, respectively. We interpret payoffs in $[0,1]$ to be in utils; that is, we assume that the cardinal utility function over outcomes is known and payoffs are stated in its units. An alternative interpretation is that there are two monetary prizes $x > y$, and $f(s) = p_s(x) \in [0,1]$ is the probability of getting the greater prize in state s.

Remark 1. *Our analysis can be easily extended to the case where, instead of $[0,1]$, the range of acts is a more general vector space. In particular, it could be formulated in the Anscombe and Aumann (1963) setting. Since our focus is on deriving the DM's subjective information structure, we abstract in the main text from deriving the utility function (which is a standard exercise) by looking directly at utility acts instead of the corresponding Anscombe-Aumann acts. Appendix A formally provides this extension and derives the utility function, thereby establishing the mapping between Anscombe-Aumann acts and utility acts.*

The interpretation of the domain is that a menu F is chosen ex-ante according to $\succeq$. This choice is made with the understanding that at the unmodeled ex-post stage, the agent will choose an act from F in the face of uncertainty about the true state of the world. The true state is realized and payoffs are received in the terminal period. The ranking $\succeq$ is governed by the anticipated choice from the menu, which, in turn, depends on what the DM expects to learn before making that choice. Subjective learning is captured in the following definition of a utility function over menus.

Definition 1. *The binary relation $\succeq$ has a **subjective-learning representation** if there is a probability measure p on $\Delta(S)$, the space of all probability measures on S, such that the function $V : \mathcal{K}(\mathcal{F}) \to \mathbb{R}$ given by*

$$V(F) = \int\limits_{\Delta(S)} \max_{f \in F} \left(\sum_{s \in S} f(s) \pi(s) \right) dp(\pi)$$

represents $\succeq$.

A subjective-learning representation has the following interpretation: before choosing an act from F, the DM can observe a signal, which allows him to update the prior probability distribution over the states. Given the posterior probabilities π, the DM chooses an act that maximizes his expected utility. Prior to observing the signal, his ex-ante expected utility is the weighted average of these maximum expected utilities, where the weights are the likelihood of receiving each signal. The probability distribution over posterior beliefs p is the information structure.[6] Figure 1 illustrates the timing of events and decisions.
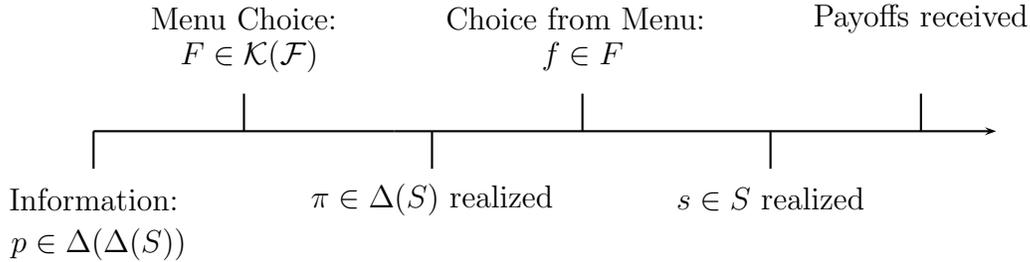
Figure 1: Subjective Learning Timeline

It is evident from the description of the model that the information structure in a subjective-learning representation is not given to us objectively. Instead, the probability measure p should be derived from choice behavior. To do this, we impose Axioms 1-6 below on $\succeq$, and then show that they are necessary and sufficient for the existence of a subjective-learning representation with uniquely identified information structure p.

2.1. Axioms and representation result

We impose six axioms on the relation $\succeq$. The first three are standard, and play the same role in this paper as they do in more familiar contexts.[7]

---

[6]Formally $p \in \Delta(\Delta(S))$, the set of all Borel probability measures over $\Delta(S)$. Note that signals in a subjective-learning representation are identified with the posterior beliefs in the support of p.

[7]We refer the reader to Dekel, Lipman, and Rustichini (2001) for a discussion of these axioms in the context of preferences over menus of lotteries.

Axiom 1 (Preference relation). *The relation $\succeq$ is complete and transitive.*

Definition 2. *Let $\alpha F + (1 - \alpha) G := \{\alpha f + (1 - \alpha) g : f \in F, \ g \in G\}$, where $\alpha f + (1 - \alpha) g$ is the act that yields $\alpha f(s) + (1 - \alpha) g(s)$ in state s.*

Axiom 2 (vNM Continuity). *If $F \succ G \succ H$ then there are $\alpha, \beta \in (0, 1)$, such that $\alpha F + (1 - \alpha) H \succ G \succ \beta F + (1 - \beta) H$.*

Axiom 3 (Nontriviality). *There are F and G such that $F \succ G$.*

Axiom 4 (Independence). *For all F, G, H, and $\alpha \in [0, 1]$,*

$$F \succeq G \Leftrightarrow \alpha F + (1 - \alpha) H \succeq \alpha G + (1 - \alpha) H.$$

In the context of choice between menus of utility acts, Axiom 4 is justified by the following two steps: First, for any two menus F, H and $\alpha \in [0, 1]$, denote by $\alpha \circ F + (1 - \alpha) \circ H$ the lottery which assigns F with probability α and H with probability $(1 - \alpha)$. The vNM independence axiom implies that for any $\alpha \in [0, 1]$, F is preferred to G if and only if $\alpha \circ F + (1 - \alpha) \circ H$ is preferred to $\alpha \circ G + (1 - \alpha) \circ H$. Second, we argue that a DM who wants to obey the vNM axioms when evaluating lotteries and acts, should be indifferent between $\alpha \circ F + (1 - \alpha) \circ H$ and $\alpha F + (1 - \alpha) H$. To see this, suppose the DM would choose f^* from F and h^* from H. Then, the realization of the lottery $\alpha \circ F + (1 - \alpha) \circ H$ together with a subsequent choice of an alternative from either F or H implies that, from the ex-ante perspective and contingent on state s, the DM faces the lottery $\alpha \circ f^*(s) + (1 - \alpha) \circ h^*(s)$. Since $f^*(s)$ and $h^*(s)$ are utils, the expected utility of this lottery is simply $\alpha f^*(s) + (1 - \alpha) h^*(s)$. The corresponding choice of $\alpha f^* + (1 - \alpha) h^*$ from $\alpha F + (1 - \alpha) H$ also yields utility $\alpha f^*(s) + (1 - \alpha) h^*(s)$, contingent on state s. Since $\alpha \circ F + (1 - \alpha) \circ H$ and $\alpha F + (1 - \alpha) H$ give the same consequence from the ex-ante perspective, the DM should be indifferent between them.

The next axiom was first proposed in Kreps (1979). It captures preference for flexibility, that is, bigger sets are weakly preferred.

Axiom 5 (Set monotonicity). *If $F \subset G$ then $G \succeq F$.*

Axiom 5 is the observable indication of information in the model; the DM likes bigger sets since more available options allow him to better adjust his choice to his updated beliefs.

The interpretation of $f(\cdot)$ as a vector of utils requires the following payoff-monotonicity axiom.

Axiom 6 (Domination). *If $f(s) \geq g(s)$ for all $s \in S$ and $f \in F$, then $F \sim F \cup \{g\}$.*

We now present our first representation theorem.

Theorem 1. *The relation $\succeq$ satisfies Axioms 1–6 if and only if it has a subjective-learning representation. Furthermore, the probability measure p is unique.*

Proof. See Appendix B.1. ∎

A related work, Dekel, Lipman, and Rustichini (2001), analyzes choice over menus of lotteries and provides a representation that suggests uncertainty about the DM's tastes (a relevant corrigendum is Dekel, Lipman, Rustichini, and Sarver (2007)). Our proof of Theorem 1 relies on a sequence of geometric arguments that establish the connection between our domain and theirs, as we now explain.

In DLR's representation,[8] the value of a menu of lotteries F is given by

$$V\left(F\right) = \int\limits_{u \in \mathcal{U}} \max_{q \in F}\left(\sum\nolimits_{z \in Z} q\left(z\right) u\left(z\right)\right) dp\left(u\right), \tag{1}$$

where Z is a finite set of prizes, q is a lottery over Z, $\mathcal{U}$ is the set of normalized Bernoulli functions u over Z, and p is a probability measure over $\mathcal{U}$ with support $\sigma\left(p\right)$. The key observation in our proof is that the set of acts on S with outcomes in $[0, 1]$ is isomorphic to the set of lotteries over $|S| + 1$ pure outcomes, where each state $s \in S$ is given the weight $\frac{f(s)}{|S|}$ and the additional state $s_{|S|+1}$ is given the weight $1 - \sum_{s \in S} \frac{f(s)}{|S|}$. Based on this observation, DLR's representation in (1) is translated in our setting to a representation of the form

$$V\left(F\right) = \frac{1}{|S|} \int\limits_{u \in \mathcal{U}} \max_{f \in F}\left(\sum\nolimits_{s \in S} f\left(s\right) u\left(s\right) + \left(|S| - \sum\nolimits_{s \in S} f\left(s\right)\right) u\left(s_{|S|+1}\right)\right) dp\left(u\right).$$

We would like to interpret each $u \in \sigma\left(p\right)$ as a probability measure over S. First, to be consistent with our notation, let $\pi\left(s\right) = u\left(s\right)$. Using Axiom 6, we show that $\pi\left(s\right) \geq \pi\left(s_{|S|+1}\right)$ for all $s \in S$ and for all $\pi \in \sigma\left(p\right)$, which means that we can normalize $\pi\left(s_{|S|+1}\right) = 0$ for all $\pi \in \sigma\left(p\right)$. The representation in Theorem 1 is then obtained by renormalizing each π to be a probability measure over S, simultaneously adjusting p to keep the relative weights across states intact.

The parameter p is uniquely identified in the representation above, because p and π are required to be probability measures. Such natural normalization does not exist in Dekel et al. (2001, 2007) and, therefore, they can only jointly identify the parameters in their representation. Unique identification underlies the behavioral comparison in Section 2.2.

---

[8]We consistently refer to the version of their representation that assumes Set Monotonicity.

2.2. More preference for flexibility and the theorem of Blackwell

Blackwell's (1951, 1953) theorem is a fundamental result in information theory. It establishes an equivalence between (i) a statistical condition under which information structure p is considered more informative than another information structure, q; and (ii) the idea that, independently of the underlying decision problem, any expected utility maximizer would prefer using p rather than q. As we pointed out above, the information structure in our paper is not objectively given but is derived from choice over menus of acts. Our goal in this section is to find the behavioral comparison (in terms of preferences over menus) which corresponds to two derived information structures that can be compared in terms of their informativeness. In particular, we use Theorem 1 to connect a notion of preference for flexibility with the DM's subjective learning. In what follows, when we discuss a particular individual i, we denote by $\succeq_i$ his preferences and by superscript i any component of his utility representation.

We first suggest a comparative notion of more preference for flexibility.

Definition 3. *DM1 has more preference for flexibility than DM2 if for all $f \in \mathcal{F}$ and for all $G \in \mathcal{K}(\mathcal{F})$,*

$$\{f\} \succeq_1 G \text{ implies } \{f\} \succeq_2 G.$$

Expressed in words, DM1 has more preference for flexibility than DM2 if whenever DM1 prefers to commit to a particular action rather than to retain an option to choose, so does DM2.[9],[10]

The next claim shows that two DMs who are comparable in terms of their preference for flexibility must agree on the ranking of singletons.

Claim 1. *Suppose DM1 has more preference for flexibility than DM2. Then $\{f\} \succeq_1 \{g\}$ if and only if $\{f\} \succeq_2 \{g\}$.*

Proof. See Appendix B.2 ■

We now compare *subjective* information structures in analogy to the notion of better information proposed by Blackwell in the context of objective information. Definition 4

---

[9]Definition 3 is analogous to the notion of "more aversion to commitment" as appears in Higashi, Hyogo, and Takeoka (2009, Definition 4.4, p. 1031) in the context of preferences over menus of lotteries.

[10]Definition 3 does not imply greater willingness to pay to add options to *any* given menu. In fact, defining more preference for flexibility this way results in an empty relation. To see this, suppose that $\succeq_1 \neq \succeq_2$ and, for simplicity, that the supports of p^1 and p^2 are finite. Using Theorem 1, there is a posterior belief π, such that $p^1(\pi) > p^2(\pi)$. It is easy to construct a menu that generates payoff $k - \delta$ under belief π and payoff k under any other belief. DM1 then would be willing to pay more than DM2 to add an act that yields payoff k under π, hence DM2 would not have more preference for flexibility than DM1. But by a symmetric argument, DM1 would also not have more preference for flexibility than DM2.

below says that an information structure is better than another one if and only if both structures induce the same prior probability distribution, and all posterior probabilities of the latter are a convex combination of the posterior probabilities of the former.

For any $p \in \Delta(\Delta(S))$, let $\sigma(p) \subset \Delta(S)$ denote the support of p.

Definition 4. *DM1 expects to be better informed than DM2 if and only if DM1's distribution of posterior beliefs is a mean-preserving spread of DM2's (in the space of probability distributions). That is, there exists a nonnegative function $k : \sigma(p^1) \times \sigma(p^2) \to \mathbb{R}_+$, satisfying $\int_{\sigma(p^1)} k(\pi, \pi') d\pi = 1$ for all $\pi' \in \sigma(p^2)$, such that*

1. $p^1(\pi) = \int_{\sigma(p^2)} k(\pi, \pi') dp^2(\pi')$ for all $\pi \in \sigma(p^1)$; and

2. $\pi'(s) = \int_{\sigma(p^1)} \pi(s) k(\pi, \pi') d\pi$ for all $\pi' \in \sigma(p^2)$ and $s \in S$.

Note that conditions (1) and (2) imply that $\int_{\sigma(p^1)} \pi(s) dp^1(\pi) = \int_{\sigma(p^2)} \pi(s) dp^2(\pi)$ for all $s \in S$, that is, the prior is the same under both p^1 and p^2.

Theorem 2. *If DM1 and DM2 have preferences that can be represented as in Theorem 1, then DM1 has more preference for flexibility than DM2 if and only if DM1 expects to be better informed than DM2.*[11]

Proof. Blackwell (1953) establishes that DM1's distribution of posterior beliefs is a mean-preserving spread of DM2's if and only if $V^1(G) \geq V^2(G)$ for any $G \in \mathcal{K}(\mathcal{F})$ (see Kihlstrom (1984) or Gollier (2001) for an illustrative proof and discussion). At the same time, $V^1(\{f\}) = V^2(\{f\})$ for any $f \in \mathcal{F}$. Hence, $V^1(\{f\}) \geq V^1(G)$ implies $V^2(\{f\}) \geq V^2(G)$. Conversely, suppose $V^2(G) > V^1(G)$ for some $G \in K(\mathcal{F})$. Then continuity implies that there exists $f \in \mathcal{F}$ with $V^2(G) > V^2(\{f\}) = V^1(\{f\}) > V^1(G)$. ∎

3. Partitional learning

We now study a more parsimonious model of learning, in which signals are deterministic, that is, they correspond to events that partition the state space. This model describes information independently of the (induced) changes in beliefs. Section 3.1 investigates the behavioral implications of this type of learning and identifies one axiom which is both necessary and sufficient to restrict the information structure in a subjective-learning representation

---

[11]The characterization of preference for flexibility via Blackwell's comparison of information structures is specific to our context, where this preference arises due to uncertainty about learning. Krishna and Sadowski (2012) provide an analogous result in a context where preference for flexibility arises due to uncertain tastes.

to produce deterministic signals. A useful application of the model is the comparison of the behavior of two DMs who learn differently without requiring them to hold the same prior beliefs, as analyzed in Section 3.2.

3.1. Partitional-learning representation

Suppose the DM knows what information he will receive contingent on the (unknown) true state of the world. He can then anticipate which act he will choose from the menu, contingent on the true state. In this case, the DM should be indifferent between choosing from the menu after learning the signal, and committing, for every state, to receive the payoff his certain choice would have generated in that state. This is the content of Axiom 7 below. In order to state the axiom, we first need to define the notion of a contingent plan from a given menu.

Definition 5. *Given $F \in \mathcal{K}(\mathcal{F})$, let $CP(F)$ be the collection of contingent plans that correspond to F, that is,*

$$CP(F) := \{g \in \mathcal{F} \mid \forall s \in S \text{ there is } f \in F \text{ with } g(s) = f(s)\}.$$

In words, a contingent plan from a menu F is an act whose outcome in any state agrees with that of some act in F. Note that $F \subseteq CP(F)$.

Axiom 7 (Contingent planning). *For every $F \in \mathcal{K}(\mathcal{F})$ there exists $f_F \in CP(F)$ such that $F \sim \{f_F\}$.*

Axiom 7 requires that any menu of acts has a corresponding contingent plan which is indifferent to it.

We now introduce the notion of a partitional-learning representation:

Definition 6. *The binary relation $\succeq$ has a **partitional-learning representation**, $(\mu, \mathcal{P})$, if (i) μ is a probability measure on S; (ii) $\mathcal{P}$ is a partition of $\sigma(\mu)$, the support of μ; and (iii) the function*

$$V(F) = \sum_{I \in \mathcal{P}} \max_{f \in F} \left(\sum_{s \in I} f(s) \mu(s) \right)$$

represents $\succeq$.

A partitional-learning representation has the following interpretation: contingent on the true state of the world being s, the DM is sure to learn a specific event I_s, with $s \in I_s \subset S$. Upon learning an event, the DM calculates his posterior beliefs by excluding all states that are not in that event and applying Bayes' law with respect to the remaining states. The DM then chooses the available act that maximizes the corresponding expected utility.

It is not difficult to see that Axiom 7 is necessary for a partitional-learning representation $(\mu, \mathcal{P})$. For any $s \in \sigma(\mu)$, let $I_s \in \mathcal{P}$ be the unique element of the partition that includes state s. For any menu F, an indifferent contingent plan $f_F \in CP(F)$ can be constructed by letting f_F pay in state s what the optimal act in F contingent on I_s pays in that state, that is,

$$f_F(s) = \left(\arg\max_{f \in F} \sum\nolimits_{s' \in I_s} f(s') \mu(s') \right)(s).$$

For example, let $S = \{s_1, s_2\}$ and $F = \{(1,0), (0.5, 0.5)\}$, which means that $CP(F) = F \cup \{(1, 0.5), (0.5, 0)\}$.[12] Let $(\mu, \mathcal{P})$ be a partitional-learning representation with $\mu(s_1) = 0.5$ and $\mathcal{P} = \{\{s_1\}, \{s_2\}\}$. Then $V(F) = V(\{(1, 0.5)\}) = 0.75$, that is, $F \sim \{(1, 0.5)\}$.

It is also easy to find a general representation as in Theorem 1 that violates Axiom 7. Consider the same S and F as above, and let $\sigma(p) = \{\pi = (0.9, 0.1), \pi' = (0.1, 0.9)\}$ with $p(\pi') = 0.5$. Then $V(F) = 0.5 \times 0.9 + 0.5 \times 0.5 = 0.7$, whereas $V(\{(0.5, 0.5)\}) = V(\{(1, 0)\}) = 0.5$, $V(\{(0, 0.5)\}) = 0.25$, and $V(\{(1, 0.5)\}) = 0.75$. This example demonstrates that Axiom 7 is easily testable; if F is finite then the set $CP(F)$ is also finite and violation of the axiom can be obtained in a finite number of observed choice situations.

Our next result establishes that such a violation of Axiom 7 can always be produced unless signals are deterministic. That is, adding Axiom 7 to the axioms in Theorem 1 is not only necessary, but also sufficient for a partitional–learning representation.

Theorem 3. *The relation $\succeq$ satisfies Axioms 1–7 if and only if it has a partitional-learning representation $(\mu, \mathcal{P})$. Furthermore, the pair $(\mu, \mathcal{P})$ is unique.*

Proof. See Appendix B.3. ∎

Theorem 3 states the conditions under which information can be uniquely identified from choice behavior as a partition of the objective state space. This partition, which is an endogenous component of the model, represents what the DM expects to know before a choice from the menu has to be made. As we point out in the introduction, a partition of the state space is a canonical formalization of information that describes signals as events in the state space.

Axiom 7 is the behavioral criterion for determining whether the analyst can indeed confine his attention to partitional learning. It implies that the objective state space S is large enough, in the sense that the DM's subjective uncertainty is measurable in S. That is, the axiom implies that the state space captures all the DM expects to know about the decision problem under consideration.

---

[12]For notational convenience, we denote here an act by an ordered pair of state contingent payoffs, $x = (x_1, x_2)$, where x_i is the payoff received in state i.

To illustrate why Axiom 7 is sufficient in Theorem 3, we first note that in the context of a subjective-learning representation as in Theorem 1, receiving deterministic signals implies that for all $\pi, \pi' \in \sigma(p)$, $\pi = \pi'$ or $\sigma(\pi) \cap \sigma(\pi') = \emptyset$. In that case we say that p is *partitional*. Roughly, if p is not partitional, we can always find a state s, two acts f and g with $0 < f(s) \neq g(s) > 0$, and a menu F that includes f and g, such that slightly changing either $f(s)$ or $g(s)$ changes $V(F)$. If $f_F \in CP(F)$ satisfies $F \sim \{f_F\}$, then either $f_F(s) \neq f(s)$ or $f_F(s) \neq g(s)$, or both. If, without loss of generality, the former holds, then $f_F \in CP(F_\varepsilon)$, but $\{f_F\} \nsim F_\varepsilon$, where F_ε is obtained from F by reducing $f(s)$ by $\varepsilon > 0$ small enough. If there happens to be $f_{F_\varepsilon} \in CP(F_\varepsilon)$ such that $F_\varepsilon \sim \{f_{F_\varepsilon}\}$, we can proceed inductively to construct a menu for which Axiom 7 is violated.

There are several advantages of the partitional-learning representation over the more general subjective-learning representation. First, as we show in Section 3.2, partitional-learning representation allows comparing individuals in terms of their preference for flexibility, even if they hold different prior beliefs. Second, partitional-learning representation can be used to forecast the DM's ex-post behavior, or choice from menus. Suppose the analyst is willing to assume that the DM's choice from a menu is governed by the subjective learning representation of Theorem 1. At the same time, however, the analyst may disagree with the DM's prior over the objective state space, as elicited from the DM's choice over menus. Suppose the DM learns by a partition of S. In that case, while behavior conditional on receiving each signal is fully determined by the DM's prior (converted into a posterior using Bayes' law), the analyst can evaluate the likelihood of each signal based on his own prior when forecasting the DM's choice frequencies. In particular, if the analyst has a degenerate prior at some state of the world (he knows the true state) then he can deduce the event the DM will learn, and consequently can predict the DM's choice from a menu with certainty. Such prediction is impossible if the DM does not face deterministic signals.

A third advantage is that partitional-learning can be used to analyze the welfare effects of a change in the informational environment. Suppose, for example, that a social planner considers a policy that would require two individuals to share their information structure. The value of such a policy is unambiguous if both individuals learn via partition, in which case the refined information structure relevant for each of them is simply the 'meet' of the two individual partitions. If, on the other hand, subjective signals are not deterministic, then assessing the value of the policy entails knowing the correlation between the individual signals; this information can not be elicited solely from an individual's choice behavior.

Lastly, the dimensionality of the parameters that describe a partitional-learning representation is significantly lower than that of the parameters used in a more general subjective-learning representation, leading to a more tractable functional form. In particular, the

measure μ in the specification of a partitional-learning representation is over the finite set S, whereas the measure p that specifies a subjective-learning representation is on the multi-dimensional vector space $\Delta(S)$.

3.2. Comparing valuations of binary bets

As shown in Section 2.2, individuals who disagree on their prior beliefs are not comparable in terms of their preference for flexibility. The partitional-learning representation facilitates a behavioral comparison of the information of individuals with different priors because partitions can be partially ranked in terms of their fineness, independent of any prior beliefs. The behavior of two individuals who expect to receive different information differs in the value they derive from the availability of binary bets, as we now describe.

For $c \in (0,1)$ and $s, s' \in S$, define the acts c^{+s} and $c^{+s}_{-s'}$ by

$$c^{+s}(\widehat{s}) = \begin{cases} 1 & \text{if } \widehat{s} = s \\ c & \text{otherwise} \end{cases} \quad \text{and} \quad c^{+s}_{-s'}(\widehat{s}) = \begin{cases} 1 & \text{if } \widehat{s} = s \\ 0 & \text{if } \widehat{s} = s' \\ c & \text{otherwise} \end{cases}.$$

Slightly abusing notation, we denote by c the constant act that always yields c.

Definition 7. *DM1 values more binary bets than DM2 if for all $s, s' \in S$ and $c \in (0,1)$,*

1. $\{c\} \sim_1 \{c^{+s}\} \Leftrightarrow \{c\} \sim_2 \{c^{+s}\}$; and

2. $\{c\} \succeq_i \{c^{+s}_{-s'}\}$ for $i = 1, 2$ and $\{c\} \succeq_1 \{c^{+s}_{-s'}, c\} \Rightarrow \{c\} \succeq_2 \{c^{+s}_{-s'}, c\}$.

Condition (1) says that the two DMs agree on whether or not payoffs in state s are valuable. Condition (2) says that having the bet $c^{+s}_{-s'}$ available is valuable to DM1 whenever it is valuable to DM2. The notion of valuing more binary bets weakens the notion of more preference for flexibility (Definition 3); Condition (2) is implied by Definition 3 and Condition (1) is implied by Claim 1.

A natural measure of the amount of information that a DM expects to receive is the fineness of his partition.

Definition 8. *The partition $\mathcal{P}$ is finer than partition $\mathcal{Q}$ if for every $I \in \mathcal{P}$ there is $I' \in \mathcal{Q}$ such that $I \subseteq I'$.*

Theorem 4. *If DM1 and DM2 have preferences that can be represented as in Theorem 3, then*

(i) DM1 values more binary bets than DM2 if and only if $\sigma\left(\mu^1\right) = \sigma\left(\mu^2\right)$ and $\mathcal{P}^1$ is finer than $\mathcal{P}^2$.

(ii) DM1 has more preference for flexibility than DM2, if and only if $\mu^1 = \mu^2$ and $\mathcal{P}^1$ is finer than $\mathcal{P}^2$.

Proof. See Appendix B.4. ∎

Theorem 4 (i) compares the behavior of two individuals who expect to learn differently, without requiring that they share the same prior beliefs; instead, the only requirement is that their prior beliefs have the same support. For example, DM1 might consider himself a better experimenter than DM2, in the sense of expecting a finer partition, even though DM2's prior is sharper. In contrast, and in line with Theorem 2, Theorem 4 (ii) states that the stronger comparison of more preference for flexibility corresponds exactly to adding the requirement that the prior beliefs are the same.

4. Subjective gradual resolution of uncertainty

In order to model a DM who expects to learn gradually over time, we now introduce the domain of *dated-menus*, where the DM can choose not only among menus but also the future time by which he will make his choice of an act from the menu. Formally, consider the domain $K\left(\mathcal{F}\right) \times [0,1]$, where a typical element (F, t) represents a menu and a time by which an alternative from the menu must be chosen. As a concrete example, consider an investor who should choose not only whether or not, but also *when* to invest, in anticipation of uncertain future profits. Delaying the investment may be costly, but it may have an option value; equipped with more information, the investor can decide to refrain from investing if he learns that profits are likely to be too low. If we denote by c_t the cost of investment at time t and by x_s the profits if state s realized, then the pair (I_t, t) corresponds to the date t option (I_t) to either invest, which pays $x_s - c_t$ in state s, or not to invest, which pays 0 in every state.

Let $\succeq^*$ be a preference relation over $\mathcal{K}\left(\mathcal{F}\right) \times [0,1]$. For each $t \in [0,1]$, define the induced binary relation $\succeq^*_t$ by $G \succeq^*_t F \Leftrightarrow (G, t) \succeq^* (F, t)$. Clearly $\succeq^*_t$ is a preference relation (it satisfies Axiom 1). We assume that each $\succeq^*_t$ also satisfies Axioms 2–6 and hence admits a subjective-learning representation, that is,

$$V_t\left(F\right) = \int_{\Delta(S)} \max_{f \in F} \left(\sum_{s \in S} f\left(s\right) \pi\left(s\right)\right) dp_t\left(\pi\right)$$

represents $\succeq^*_t$, and $p_t\left(\cdot\right)$ is unique.

For each $F \in \mathcal{K}(\mathcal{F})$, define the induced binary relation $\succeq_F^*$ by $t \succeq_F^* t' \Leftrightarrow (F,t) \succeq^* (F,t')$. Again, $\succeq_F^*$ satisfies Axiom 1. We impose on $\succeq_F^*$ the following additional axioms:

Axiom 8 (Preference for choosing late). *For all $t, t' \in [0,1]$ and $F \in \mathcal{K}(\mathcal{F})$,*

$$t \geq t' \Rightarrow t \succeq_F^* t'.$$

If the DM expects uncertainty to resolve over time, then waiting enables him to make a more informed choice from the menu. The next axiom rules out intrinsic preference for the timing of choice, independently of the instrumental value of information.

Axiom 9 (Time-independent singleton preferences). *For all $t, t' \in [0,1]$ and $f \in \mathcal{F}$,*

$$t \sim_{\{f\}}^* t'.$$

Since the singleton menu $\{f\}$ does not leave the DM any flexibility to adjust his choice to new information, the time component should play no role in its evaluation.

Definition 9. *The collection of measures $(p_t)_{t \in [0,1]}$ is a **gradual-learning representation** if the function V^* given by*

$$V^*(F,t) = \int_{\Delta(S)} \max_{f \in F} \left(\sum_{s \in S} f(s) \pi(s) \right) dp_t(\pi)$$

represents $\succeq^*$, and p_t is a mean preserving spread of $p_{t'}$, whenever $t \geq t'$.

Theorem 5. *Let $\succeq^*$ be a preference relation over $\mathcal{F} \times [0,1]$. The following two statements are equivalent:*

(i) For each $t \in [0,1]$, the relation $\succeq_t^*$ satisfies Axioms 2–6, and for each $F \in \mathcal{K}(\mathcal{F})$, the relation $\succeq_F^*$ satisfies Axiom 8 and Axiom 9.

(ii) The relation $\succeq^*$ has a gradual-learning representation. Furthermore, the collection $(p_t)_{t \in [0,1]}$ is unique.

Proof. See Appendix B.5. ■

The key idea behind the proof of Theorem 5 is the observation that Axioms 8 and 9 imply that $\succeq_t^*$ has more preference for flexibility than $\succeq_{t'}^*$ whenever $t > t'$. The result then follows by applying Theorem 2.

A DM who expects to learn gradually over time may face an intertemporal trade-off (as in the example in the beginning of this section). If he anticipates late resolution of uncertainty, he would like to postpone his decision until the uncertainty is resolved. But waiting might

18

be costly, in the sense that the set of available options at that later time will be inferior to the one available earlier. Theorem 5 suggests a way to resolve this intertemporal trade-off; it pins down how the DM's knowledge will be improved through time and how this improved knowledge affects the values of different choice problems (menus).

Theorem 5 can be readily specialized to the case where for each $t \in [0,1]$, $\succeq_t^*$ satisfies Axioms 2–7 and hence admits a partitional-learning representation (as in Section 3). In that case, it captures a DM who has in mind a filtration.

Definition 10. *The pair* $\left(\mu, \{\mathcal{P}_t\}_{t \in [0,1]}\right)$ *is a **learning by filtration representation** if (i)* μ *is a probability measure on S; (ii)* $\{\mathcal{P}_t\}_{t \in [0,1]}$ *is a filtration on $\sigma(\mu)$ indexed by $t \in [0,1]$;[13] and*

$$V^*(F,t) = \sum_{I \in \mathcal{P}_t} \max_{f \in F} \left(\sum_{s \in I} f(s)\, \mu(s)\right)$$

represents $\succeq^*$.

Corollary 1. *Let $\succeq^*$ be a preference relation over $\mathcal{F} \times [0,1]$. The following two statements are equivalent:*

(i) For each $t \in [0,1]$, the relation $\succeq_t^*$ satisfies Axioms 2–7, and for each $F \in \mathcal{K}(\mathcal{F})$, the relation $\succeq_F^*$ satisfies Axiom 8 and Axiom 9.

(ii) The relation $\succeq$ has a learning by filtration representation. Furthermore, the pair $\left(\mu, \{\mathcal{P}_t\}_{t \in [0,1]}\right)$ is unique.

Proof. It is sufficient to observe that under the additional Axiom 7, p_t is a mean preserving spread of $p_{t'}$ if and only if P_t is finer than $P_{t'}$ and $\mu_t = \mu_{t'}$. ■

Under the assumptions underlying the learning by filtration representation, an immediate implication of Theorem 4 (ii) is that DM1 has more preference for flexibility than DM2 for all t, if and only if $\mu^1 = \mu^2$ and he has a finer filtration, that is, the partition P_t^1 is finer than P_t^2 for all t. Intuitively, DM1 has more preference for flexibility than DM2 for all t if and only if he expects to learn earlier. Reconsider the investment example outlined in the beginning of the section, where a menu (I_t, t) corresponds to the date t option to either invest, which pays $x_s - c_t$ in state s, or not to invest, which pays 0 in every state. In particular, the menu $(I_0, 0)$ corresponds to the option of investing immediately, before any new information can arrive. Suppose that (i) $c_0 = 0$ and that $t > t'$ implies $c_t > c_{t'}$; (ii) both DMs would invest given their initial knowledge, that is, they are indifferent between $(I_0, 0)$ and the act that pays x_s in state s; and (iii) the decision to invest can be delayed by some arbitrary but fixed

---

[13]Slightly abusing notation, we identify a filtration with a right-continuous and nondecreasing (with respect to set inclusion) function from $[0,1]$ to $2^{\sigma(\mu)}$.

amount of time t. In this case, if DM1 learns earlier than DM2, then if DM1 does not value the option to delay the decision, neither does DM2.

Appendices

A. Anscombe and Aumann setting

Since our focus is on deriving the DM's subjective information structure, we abstract in the body of the paper from deriving the utility function by looking directly at utility acts. In this appendix, we show that our analysis can be easily extended to the case where the range of acts is the set of lotteries, known as the Anscombe and Aumann (1963) setting.

Let S be a finite state space, and X be a compact metric space. The set of all Borel probability measures, called lotteries, on X is denoted by $\Delta(X)$, which is endowed with the weak convergence topology. An Anscombe-Aumann act (AA-act) is a function $f : S \to \Delta(X)$. Let $\mathcal{F}^*$ be the set of all AA-acts. We consider preference $\succeq$ over the set $\mathcal{K}(\mathcal{F}^*)$ of all menus of AA-acts. A mixture operation between two menus is defined by the mixture of two AA-acts as usual.

Axioms 1, 2, 4, and 5 can be directly translated into the current setting. To reformulate Axiom 6, for all $f, g \in \mathcal{F}^*$, we say that f dominates g if $\{f(s)\} \succeq \{g(s)\}$ for all $s \in S$. That is, a dominant act f gives the DM a better lottery in terms of commitment ranking in every state of the world. The following is an analogue of Axiom 6:

Axiom 10 (Domination*). *If f dominates g and $f \in F$, then $F \sim F \cup \{g\}$.*

Since a dominated act always gives less utilities no matter what information arrives in the future, it is never chosen over a dominant act. Therefore, the DM does not care about adding a dominated act into a menu, as described in the axiom.

The appropriate version of Axiom 3 is

Axiom 11 (Nontriviality*). *there exist lotteries $\ell, \ell' \in \Delta(X)$ with $\{\ell\} \succ \{\ell'\}$.*

Now suppose that preference $\succeq$ over $\mathcal{K}(\mathcal{F}^*)$ satisfies Axioms 1, 2, 4, 5, 10, and 11. By Axioms 1, 2, 11, and 4, there exists a non-constant mixture linear function $u : \Delta(X) \to [0, 1]$ representing the commitment preference over $\Delta(X)$. By using this utility function as a measuring rod, for any $f \in \mathcal{F}^*$, we can define the corresponding utility act $u(f) : S \to [0, 1]$ by $u(f)(s) := u(f(s))$ for all s. Furthermore, for all $F \in \mathcal{K}(\mathcal{F}^*)$, define the corresponding menu in $\mathcal{K}(\mathcal{F})$ by $F_u := \{u(f) \in \mathcal{F} \mid f \in F\}$. The induced preference over $\mathcal{K}(\mathcal{F})$ is defined as $F_u \succeq G_u$ if and only if $F \succeq G$.

The following proposition states that the Anscombe-Aumann setting can be reduced to our setting of utility acts. A proof is omitted.

Proposition 1. *If Preference $\succeq$ over $\mathcal{K}(\mathcal{F}^*)$ satisfies Axioms 1, 2, 4, 5, 10, and 11, then the induced preference over $\mathcal{K}(\mathcal{F})$ satisfies Axioms 1–6.*

In order to conduct interpersonal comparisons about information structures, it suffices to assume that two individuals $\succeq_i$, $i = 1, 2$, have the same risk preference, that is, for all $\ell, \ell' \in \Delta(X)$, $\{\ell\} \succeq_1 \{\ell'\}$ if and only if $\{\ell\} \succeq_2 \{\ell'\}$. Then, without loss of generality, we can assume that their risk preferences are represented by the same expected utility function u over $\Delta(X)$, and adopt the same interpersonal comparisons as in the previous sections.

B. Proofs

B.1. Proof of Theorem 1

It is easily verified that any binary relation $\succeq$ with a subjective-learning representation satisfies the axioms. We proceed to show the sufficiency of the axioms.

We can identify $\mathcal{F}$ with the set of all $k-$dimensional vectors, where each entry is in $[0, 1]$. For reasons that will become clear below, we now introduce an artificial state, s_{k+1}. Let

$$\mathcal{F}' := \left\{ f' \in [0, 1]^k \times [0, k] \,\middle|\, \sum_{i=1}^{k+1} f'(s_i) = k \right\}.$$

Note that the $k + 1$ component in f' equals $k - \sum_{i=1}^k f'(s_i)$. For $f' \in \mathcal{F}'$, denote by $f'^k \in \mathcal{F}$ the vector that agrees with the first k components of f'. Since $\mathcal{F}$ and $\mathcal{F}'$ are isomorphic, we can look at a preference relation on $\mathcal{K}(\mathcal{F}')$, $\succeq_*$, defined by: $F' \succeq_* G' \Leftrightarrow F \succeq G$, where $F := \left\{ f \in \mathcal{F} \,\middle|\, f = f'^k \text{ for some } f' \in F' \right\}$ and analogously for G.

Claim 2. *The relation $\succeq_*$ satisfies the independence axiom.*

Proof. Using the definition of $\succeq_*$ and Axiom 4, we have, for all F', G', and H' in $\mathcal{K}(\mathcal{F}')$ and for all $\alpha \in [0, 1]$,

$$F' \succeq_* G' \Leftrightarrow F \succeq G \Leftrightarrow \alpha F + (1 - \alpha) H \succeq \alpha G + (1 - \alpha) H \Leftrightarrow$$

$$(\alpha F + (1 - \alpha) H)' \succeq_* (\alpha G + (1 - \alpha) H)' \Leftrightarrow \alpha F' + (1 - \alpha) H' \succeq_* \alpha G' + (1 - \alpha) H'.$$

∎

Let

$$\mathcal{F}'' := \left\{ f' \in [0, k]^{k+1} \,\middle|\, \sum_{i=1}^{k+1} f'(s_i) = k \right\}.$$

21

Let $F^{k+1} := \left\{\left(\frac{k}{k+1}, ..., \frac{k}{k+1}\right)\right\} \in \mathcal{K}\left(\mathcal{F}'\right)$. Observe that for $F'' \in \mathcal{F}''$ and $\varepsilon < \frac{1}{k^2}$, $\varepsilon F'' + (1 - \varepsilon) F^{k+1} \in \mathcal{K}\left(\mathcal{F}'\right)$. Define $\succeq_{**}$ on $\mathcal{K}\left(\mathcal{F}''\right)$ by $F'' \succeq_{**} G'' \Leftrightarrow \varepsilon F'' + (1 - \varepsilon) F^{k+1} \succeq_* \varepsilon G'' + (1 - \varepsilon) F^{k+1}$ for all $\varepsilon < \frac{1}{k^2}$.

Claim 3. *The relation $\succeq_{**}$ is the unique extension of $\succeq_*$ to $\mathcal{K}\left(\mathcal{F}''\right)$ that satisfies the independence axiom.*

Proof. Note that the $(k+1)$-dimensional vector $\left(\frac{k}{k+1}, ..., \frac{k}{k+1}\right) \in int\mathcal{F}' \subset \mathcal{F}''$, hence $F^{k+1} \subset int\mathcal{F}' \subset \mathcal{F}''$. We now show that $\succeq_{**}$ satisfies independence. For any $F'', G'', H'' \in \mathcal{K}\left(\mathcal{F}''\right)$ and $\alpha \in [0, 1]$,

$$F'' \succeq_{**} G'' \Leftrightarrow \varepsilon F'' + (1 - \varepsilon) F^{k+1} \succeq_* \varepsilon G'' + (1 - \varepsilon) F^{k+1} \Leftrightarrow$$
$$\alpha \left(\varepsilon F'' + (1 - \varepsilon) F^{k+1}\right) + (1 - \alpha) \left(\varepsilon H'' + (1 - \varepsilon) F^{k+1}\right)$$
$$= \varepsilon \left(\alpha F'' + (1 - \alpha) H''\right) + (1 - \varepsilon) F^{k+1} \succeq_*$$
$$\alpha \left(\varepsilon G'' + (1 - \varepsilon) F^{k+1}\right) + (1 - \alpha) \left(\varepsilon H'' + (1 - \varepsilon) F^{k+1}\right)$$
$$= \varepsilon \left(\alpha G'' + (1 - \alpha) H''\right) + (1 - \varepsilon) F^{k+1} \Leftrightarrow \alpha F'' + (1 - \alpha) H'' \succeq_{**} \alpha G'' + (1 - \alpha) H''.$$

The first and third $\Leftrightarrow$ is by the definition of $\succeq_{**}$. The second $\Leftrightarrow$ is by Claim 2.[14] This argument shows that a linear extension exists. To show uniqueness, let $\widehat{\succeq}$ be any linear extension over $\mathcal{K}\left(\mathcal{F}''\right)$ of $\succeq$. By independence, $F'' \widehat{\succeq} G'' \Leftrightarrow \varepsilon F'' + (1 - \varepsilon) F^{k+1} \widehat{\succeq} \varepsilon G'' + (1 - \varepsilon) F^{k+1}$. Since $\widehat{\succeq}$ extends $\succeq_*$, they must agree on $\mathcal{K}\left(\mathcal{F}'\right)$. Therefore,

$$\varepsilon F'' + (1 - \varepsilon) F^{k+1} \widehat{\succeq} \varepsilon G'' + (1 - \varepsilon) F^{k+1} \Leftrightarrow \varepsilon F'' + (1 - \varepsilon) F^{k+1} \succeq_* \varepsilon G'' + (1 - \varepsilon) F^{k+1}.$$

By combining the two equivalences above, we conclude that defining $\widehat{\succeq}$ by $F'' \widehat{\succeq} G'' \Leftrightarrow \varepsilon F'' + (1 - \varepsilon) F^{k+1} \succeq_* \varepsilon G'' + (1 - \varepsilon) F^{k+1}$ is the only admissible extension of $\succeq_*$. ∎

The domain $\mathcal{K}\left(\mathcal{F}''\right)$ with the Hausdorff metric is formally equivalent to that of Dekel, Lipman, Rustichini, and Sarver (2007, henceforth DLRS) with $k + 1$ prizes. (The unit simplex is obtained by rescaling all elements of $\mathcal{F}''$ by $1/k$, that is, by redefining $\mathcal{F}''$ as $\left\{f' \in [0, 1]^{k+1} : \sum_{i=1}^{k+1} f'(s_i) = 1\right\}$.) Applying Theorem 2 in DLRS,[15] one obtains the follow-

---

[14] The $(=)$ sign in the third and in fifth lines are due to the fact that F^{k+1} is a singleton menu. For a singleton menu $\{f\}$ and $\alpha \in (0, 1)$,

$$\alpha \{f\} + (1 - \alpha) \{f\} = \{f\}$$

while, for example,

$$\alpha \{f, g\} + (1 - \alpha) \{f, g\} = \{f, g, \alpha f + (1 - \alpha) g, \alpha g + (1 - \alpha) f\},$$

is not generally equal to $\{f, g\}$.

[15] DLRS provide a supplemental appendix which shows that, for the purpose of the theorem, their stronger continuity assumption can be relaxed to the weaker notion of vNM continuity used in the present paper.

ing representation of $\succeq_{**}$:

$$\widehat{V}(F'') = \int_{\mathcal{M}(S)} \max_{f'' \in F''} \left(\sum_{s \in S \cup \{s_{k+1}\}} f''(s) \widehat{\pi}(s) \right) d\widehat{p}(\widehat{\pi}),$$

where $\mathcal{M}(S) := \left\{ \widehat{\pi} \,\middle|\, \sum_{s \in S \cup \{s_{k+1}\}} \widehat{\pi}(s) = 0 \text{ and } \sum_{s \in S \cup \{s_{k+1}\}} (\widehat{\pi}(s))^2 = 1 \right\}$. Given the normalization of $\widehat{\pi} \in \mathcal{M}(S)$, $\widehat{p}(\cdot)$ is a unique probability measure. Note that $\widehat{V}$ also represents $\succeq_*$ when restricted to its domain, $\mathcal{K}(\mathcal{F}')$.

We aim for a representation of $\succeq$ of the form

$$V(F) = \int_{\Delta(S)} \max_{f \in F} \left(\sum_{s \in S} f(s) \pi(s) \right) dp(\pi),$$

where $f(\cdot)$ is a vector of utils and $p(\cdot)$ is a unique probability measure on $\Delta(S)$, the space of all probability measures on S.

We now explore the additional constraint imposed on $\widehat{V}$ by Axiom 6 and the definition of $\succeq_*$.

Claim 4. $\widehat{\pi}(s_{k+1}) \leq \widehat{\pi}(s)$ for all $s \in S$, $\widehat{p}$−almost surely.

Proof. Suppose there exists some event $E \subset \mathcal{M}(S)$ with $\widehat{p}(E) > 0$ and $\widehat{\pi}(s_{k+1}) > \widehat{\pi}(s)$ for some $s \in S$ and all $\widehat{\pi} \in E$. Let $f' = (0, 0, ..., 0, \varepsilon, 0, ..., k - \varepsilon)$, where ε is received in state s and $k - \varepsilon$ is received in state s_{k+1}. Let $g' = (0, 0, ..0, 0, 0, ..., k)$. Then $\{f', g'\} \succ_* \{f'\}$. Take $F' = \{f'\}$ (so that $F' \cup \{g'\} \succ_* F'$). But note that Axiom 6 and the definition of $\succeq_*$ imply that $F' \sim_* F' \cup \{g'\}$, which is a contradiction. ■

Given our construction of $\widehat{V}$, there are two natural normalizations that allow us to replace the measure $\widehat{p}$ on $\mathcal{M}(S)$ with a unique probability measure p on $\Delta(S)$.

First, since s_{k+1} is an artificial state, the representation should satisfy $\pi(s_{k+1}) = 0$, p−almost surely. For all $s \in S$ and for all $\widehat{\pi}$, define $\xi(\widehat{\pi}(s)) := \widehat{\pi}(s) - \widehat{\pi}(s_{k+1})$. Since $\sum_{i=1}^{k+1} f'(s_i) = k$ and ξ simply adds a constant to every $\widehat{\pi}$,

$$\underset{f'' \in F''}{\arg\max} \left(\sum_{s \in S \cup \{s_{k+1}\}} f''(s) \xi(\widehat{\pi}(s)) \right) = \underset{f'' \in F''}{\arg\max} \left(\sum_{s \in S \cup \{s_{k+1}\}} f''(s) \widehat{\pi}(s) \right)$$

for all $\widehat{\pi} \in \sigma(\widehat{p})$. Furthermore, by Claim 4, $\xi(\widehat{\pi}(s)) \geq 0$ for all $s \in S$, $\widehat{p}$−almost surely.

Second, we would like to transform $\xi \circ \widehat{\pi}$ into a probability measure π. Let

$$\pi(s) := \xi(\widehat{\pi}(s)) / \left(\sum_{s' \in S} \xi(\widehat{\pi}(s')) \right).$$

(recall that $\xi(\widehat{\pi}(s_{k+1})) = 0$). Since this transformation affects the relative weight given to

event $E \subset \mathcal{M}(S)$ in the representation, we need p to be a probability measure on E that offsets this effect. The identification result in DLRS implies that this p is unique and can be calculated via the Radon-Nikodym derivative

$$\frac{dp(\pi)}{d\widehat{p}(\widehat{\pi})} = \frac{\sum_{s \in S} \xi(\widehat{\pi}(s))}{\int_{\mathcal{M}(S)} \left(\sum_{s \in S} \xi(\widehat{\pi}(s))\right) d\widehat{p}(\widehat{\pi})}.$$

Therefore, $\succeq$ is represented by

$$V(F) = \int_{\Delta(S)} \max_{f \in F} \left(\sum_{s \in S} f(s)\pi(s)\right) dp(\pi),$$

and the measure p is unique.

B.2. Proof of Claim 1

Let $G = \{g\}$ for some $g \in \mathcal{F}$. Applying Definition 3 implies that if $\{f\} \sim_1 \{g\}$ then $\{f\} \sim_2 \{g\}$. That is, any indifference set of the restriction of $\succeq_1$ to singletons is a subset of some indifference set of the restriction of $\succeq_2$ to singletons. The linearity (in probabilities) of the restriction of $V^i(\cdot)$ to singletons implies that these indifference sets are planes that separate any n–dimensional unit simplex, for $n \leq (|S| - 1)$. Therefore, the indifference sets of the restriction of $\succeq_1$ and $\succeq_2$ to singletons must coincide. Since the restrictions of $\succeq_1$ and of $\succeq_2$ to singletons share the same indifference sets and since both relations are monotone, they must agree on all upper and lower contour sets. In particular, $\{f\} \succeq_1 \{g\}$ if and only if $\{f\} \succeq_2 \{g\}$.

B.3. Proof of Theorem 3

First observe that any partitional-learning representation $(\mu, \mathcal{P})$ (Definition 6) can be written as a subjective-learning representation (Definition 1), where the information structure p has the property that $\pi = \pi'$ or $\sigma(\pi) \cap \sigma(\pi') = \emptyset$ for all $\pi, \pi' \in \sigma(p)$. In this case we say that p is partitional. If p is partitional, then for all $s \in S$ with $\int_{\Delta(S)} \pi(s) dp(\pi) > 0$, the prior probability μ is defined by $\mu(s) := p(\pi_s)\pi_s(s)$, where $\pi_s \in \Delta(S)$ denotes the unique posterior such that $s \in \sigma(\pi_s)$.

For any $F \in \mathcal{K}(\mathcal{F})$, let $CP(F)$ be the collection of contingent plans that correspond to F, that is,

$$CP(F) := \{g \in \mathcal{F} \mid \forall s \in S \text{ there is } f \in F \text{ with } g(s) = f(s)\}.$$

24

Let $ICP(F)$ be the contingent plans that are indifferent to F,

$$ICP(F) := \{g \in CP(F) \mid \{g\} \sim F\}.$$

We now show that any partitional-learning representation satisfies Axiom 7.

Claim 5. *If p is partitional, then for every $F \in \mathcal{K}(\mathcal{F})$, $ICP(F) \neq \varnothing$.*

Proof. Suppose that p is partitional. Using Theorem 1, for any set F we have

$$V(F) = \sum_{\pi} \max_{f \in F} \left(\sum_s f(s) \pi(s)\right) p(\pi)$$
$$= \sum_s \max_{f \in F} \left(\sum_{s'} f(s') \pi_s(s')\right) \widetilde{p}(s),$$

where $\widetilde{p}$ is a probability distribution over S such that for all $\pi \in \sigma(p)$, $\sum_{s \in \sigma(\pi)} \widetilde{p}(s) = p(\pi)$.
Let

$$f_s = \arg\max_{f \in F} \left(\sum_{s'} f(s') \pi_s(s')\right)$$

and let f_F be an act such that for all s,

$$f_F(s) = f_s(s).$$

Clearly $f_F \in CP(F)$. We have,

$$V(\{f_F\}) = \sum_s \sum_{s'} f_F(s') \pi_s(s') \widetilde{p}(s)$$
$$= \sum_{s'} f_F(s') \sum_s \pi_s(s') \widetilde{p}(s)$$
$$= \sum_{s'} f_{s'}(s') \sum_{\pi} \pi(s') p(\pi)$$
$$= \sum_{s'} f_{s'}(s') \mu(s')$$

and

$$V\left(F\right) = \sum_{s} \max_{f \in F}\left(\sum_{s'} f\left(s'\right)\pi_{s}\left(s'\right)\right)\widetilde{p}\left(s\right)$$

$$= \sum_{s}\sum_{s'} f_{s}\left(s'\right)\pi_{s}\left(s'\right)\widetilde{p}\left(s\right)$$

$$= \sum_{s'}\sum_{s} f_{s}\left(s'\right)\pi_{s}\left(s'\right)\widetilde{p}\left(s\right)$$

$$= \sum_{s'}\sum_{s \in \sigma\left(\pi_{s'}\right)} f_{s}\left(s'\right)\pi_{s}\left(s'\right)\widetilde{p}\left(s\right) \ \ \left(\text{since } \pi_{s'}\left(s\right) = 0 \text{ if } s \notin \sigma\left(\pi_{s'}\right)\right)$$

$$= \sum_{s'}\sum_{s \in \sigma\left(\pi_{s'}\right)} f_{s'}\left(s'\right)\pi_{s}\left(s'\right)\widetilde{p}\left(s\right) \ \ \left(f_{s} = f_{s'} \text{ since } \pi_{s} = \pi_{s'} \text{ if } s \in \sigma\left(\pi_{s'}\right)\right)$$

$$= \sum_{s'} f_{s'}\left(s'\right)\sum_{s \in \sigma\left(\pi_{s'}\right)}\pi_{s}\left(s'\right)\widetilde{p}\left(s\right)$$

$$= \sum_{s'} f_{s'}\left(s'\right)\mu\left(s'\right) \ \ \left(\text{since } \pi_{s'}\left(s\right) = 0 \text{ if } s \notin \sigma\left(\pi_{s'}\right)\right).$$

Therefore $V\left(F\right) = \sum_{s'} f_{s'}\left(s'\right)\mu\left(s'\right) = V\left(\{f_{F}\}\right)$, which means that $f_{F} \in ICP\left(F\right)$. ∎

The next two claims establish that Axiom 7 implies a partitional-learning representation. We first show that Axiom 7 implies that the cardinality of the support of p, $\sigma\left(p\right)$, is bounded above by the number of states.

Claim 6. $|\sigma\left(p\right)| \leq |S|$.

Proof. Suppose $|\sigma\left(p\right)| > |S|$. Let $\{\pi_{i} \in \sigma\left(p\right)\}_{i=1,..,|S|+1}$ be a set of $|S| + 1$ arbitrary beliefs. For small enough $\varepsilon > 0$, there is a collection of acts $\{f_{0}^{i}\}_{i=1,..,|S|+1}$ such that for all i and $j \neq i$, $\pi \in N_{\varepsilon}\left(\pi_{i}\right)$, the ε neighborhood of π_{i}, implies

$$\sum_{s} f_{0}^{i}\left(s\right)\pi\left(s\right) > \sum_{s} f_{0}^{j}\left(s\right)\pi\left(s\right).$$

Because of the strict inequality, we can assume that $f_{0}^{j}\left(s\right) \neq f_{0}^{i}\left(s\right)$ for any $j \neq i$ and all $s \in S$. Let $F_{0} = \{f_{0}^{i}\}_{i=1,..,|S|+1}$. Then $|CP\left(F_{0}\right)| = \left(|S| + 1\right)^{|S|} < \infty$. If $ICP\left(F_{0}\right) = \varnothing$ we immediately have a violation of Axiom 7. Suppose instead that $|ICP\left(F_{0}\right)| = k > 0$.

Proceed inductively as follows: Suppose $F_{\left(l-1\right)} = \left\{f_{\left(l-1\right)}^{i}\right\}_{i=1,..,|S|+1}$ where $f_{\left(l-1\right)}^{j}\left(s\right) \neq f_{\left(l-1\right)}^{i}\left(s\right)$ for any $j \neq i$ and all $s \in S$. Further, suppose that $\left|ICP\left(F_{\left(l-1\right)}\right)\right| \leq k - l + 1$, and that there is $\varepsilon > 0$, such that for all i and $j \neq i$, $\pi \in N_{\varepsilon}\left(\pi_{i}\right)$ implies

$$\sum_{s} f_{\left(l-1\right)}^{i}\left(s\right)\pi\left(s\right) > \sum_{s} f_{\left(l-1\right)}^{j}\left(s\right)\pi\left(s\right). \tag{2}$$

Arbitrarily pick an act $g \in ICP\left(F_{\left(l-1\right)}\right)$. Since $\left|F_{\left(l-1\right)}\right| > |S|$, there exists $i^{*} \in \{1,..,|S|+1\}$ such that $g\left(s\right) \neq f_{\left(l-1\right)}^{i^{*}}\left(s\right)$ for any s. For any $\delta > 0$ and an act f, let $f - \delta$ be the act such that

$$\left(f - \delta\right)\left(s\right) = \max\left\{f\left(s\right) - \delta, 0\right\}.$$

Let

$$
f_l^i := \begin{cases} f_{(l-1)}^i & \text{if } i \neq i^* \\ f_{(l-1)}^i - \varepsilon & \text{if } i = i^* \end{cases}
$$

and $F_l := \{f_l^i\}_{i=1,..,|S|+1}$. By Equation (2) $f_{(l-1)}^i$ is a maximizer in $F_{(l-1)}$ under π_i, and hence

$$
V(F_l) < V(F_{(l-1)}) = V(\{g\}).
$$

For $\varepsilon > 0$ small enough and for all i and $j \neq i$, $f_l^j(s) \neq f_l^i(s)$ for all $s \in S$, and $\pi \in N_\varepsilon(\pi_i)$ still imply

$$
\sum_s f_l^i(s)\,\pi(s) > \sum_s f_l^j(s)\,\pi(s).
$$

For any act $h \in CP(F_{(l-1)})$, let

$$
h'(s) = \begin{cases} h(s) & \text{if } h(s) \neq f_{(l-1)}^{i^*}(s) \\ h(s) - \varepsilon & \text{if } h(s) = f_{(l-1)}^{i^*}(s) \end{cases}.
$$

Note that $h' \in CP(F_l)$ if and only if $h \in CP(F_{(l-1)})$. In particular, $g \in CP(F_l)$. Furthermore, $g \notin ICP(F_l)$ and, for $\varepsilon > 0$ small enough, $h' \notin ICP(F_l)$ if $h \notin ICP(F_{(l-1)})$. Therefore, $|ICP(F_l)| \leq k - l$. Terminate the induction whenever $ICP(F_l) = \varnothing$, which is a violation of Axiom 7. Note that the induction will terminate in at most k steps. ■

For proving the next claim, we need to use the notion of saturated menus (initially introduced in Dillenberger and Sadowski (2012a, Definition 4)), that we now describe.

Given $f \in \mathcal{F}$, let f_s^x be the act

$$
f_s^x(s') = \begin{cases} f(s') & \text{if } s' \neq s \\ x & \text{if } s' = s \end{cases}.
$$

Let $\sigma(f) := \{s \in S \,|\, f(s) > 0\} = \{s \in S \,|\, f_s^0 \neq f\}$.

Definition 11. A menu $F \in \mathcal{K}(\mathcal{F})$ is saturated if it satisfies

(i) for all $f \in F$ and for all $s \in \sigma(f)$, $F \succ (F \setminus \{f\}) \cup \{f_s^0\}$;

(ii) for all $f \in F$ and $s \notin \sigma(f)$, there exists $\overline{\varepsilon} > 0$ such that $F \sim F \cup f_s^\varepsilon$ for all $\varepsilon < \overline{\varepsilon}$; and

(iii) if $G \nsubseteq F$ then $F \cup G \sim (F \cup G) \setminus \{g\}$ for some $g \in F \cup G$.

Claim 7. If $ICP(F) \neq \varnothing$ for every $F \in \mathcal{K}(\mathcal{F})$, then p is partitional.

Proof. We prove the contrapositive of Claim 7, that is, we show that if p is not partitional, then there exists $F \in \mathcal{K}(\mathcal{F})$ such that $ICP(F) = \varnothing$. Claim 6 implies that $\sigma(p)$ is finite

27

and hence, by Claim 1 in Dillenberger and Sadowski (2012a), a saturated menu exists. Dillenberger and Sadowski (2012a, Claim 2) show that every element in a saturated menu is a unique maximizer for a belief in $\sigma(p)$. Therefore, there is a saturated menu F_0 for which $f^1, f^2 \in F_0$ implies that either $f^1(s') \neq f^2(s')$ or $f^1(s') = f^2(s') = 0$ for all $s' \in S$. If $ICP(F_0) = \varnothing$ we immediately have a violation of Axiom 7. Assume then that $|ICP(F_0)| = k > 0$.

Suppose p is not partitional. Then there are $\pi, \pi' \in \sigma(p)$, $\pi \neq \pi'$, and a state s such that $s \in \sigma(\pi) \cap \sigma(\pi')$. This implies that in any saturated menu F, there are acts f^1 and f^2 such that for $i = 1, 2$, $f^i(s) > 0$ and $F \succ (F \backslash \{f^i\}) \cup \left\{ f_s^{if^i(s)-\varepsilon} \right\}$ for any $\varepsilon > 0$ (see Definition 11).

Proceed inductively as follows: suppose $F_{(l-1)}$ is a saturated menu such that (i) $f^1, f^2 \in F_{(l-1)}$ implies that either $f^1(s') \neq f^2(s')$ or $f^1(s') = f^2(s') = 0$ for all $s' \in S$; and (ii) $\left| ICP\left(F_{(l-1)}\right) \right| \leq k - l + 1$. Pick any $g \in ICP\left(F_{(l-1)}\right)$ and note that there is $f \in F_{(l-1)}$ with $g(s) \neq f(s) > 0$. Let $F_l := \left(F_{(l-1)} \backslash \{f\}\right) \cup \left\{ f_s^{f(s)-\varepsilon} \right\}$. For $\varepsilon > 0$ small enough, F_l is saturated, and $f^1, f^2 \in F_l$ implies $f^1(s') \neq f^2(s')$ or $f^1(s') = f^2(s') = 0$ for all $s' \in S$. For any act $h \in CP\left(F_{(l-1)}\right)$, let

$$
h'(s') = \begin{cases} h(s') - \varepsilon & \text{if } s' = s \text{ and } h(s) = f(s) \\ h(s') & \text{otherwise} \end{cases}.
$$

Note that $h' \in CP(F_l)$ if and only if $h \in CP\left(F_{(l-1)}\right)$. In particular, $g \in CP(F_l)$. For $\varepsilon > 0$ small enough, $h' \notin ICP(F_l)$ if $h \notin ICP\left(F_{(l-1)}\right)$ and $g \notin ICP(F_l)$, as $\{g\} \sim F_{(l-1)} \succ F_l$. Therefore, $|ICP(F_l)| \leq k - l$. Terminate the induction whenever $ICP(F_l) = \varnothing$, which is a violation of Axiom 7. Note that the induction will terminate in at most k steps. ∎

This concludes the proof of Theorem 3.

B.4. Proof of Theorem 4

(i) Let $\succeq$ be represented as in Theorem 3 and the acts c^{+s} and $c_{-s'}^{+s}$ as defined in the text. We make the following two observations: first, $\{c^{+s}\} \sim \{c\}$ if and only if $s \notin \sigma(\mu)$. Second, using the property that conditional on any $I \ni s, s'$,

$$
\frac{\Pr(s \mid I)}{\Pr(s' \mid I)} = \frac{\mu(s)}{\mu(s')},
$$

$\{c\} \succ \{c_{-s'}^{+s}\}$ implies that $\sum_{\widehat{s} \in I} c_{-s'}^{+s}(\widehat{s}) \, \mu(\widehat{s}) > c \sum_{\widehat{s} \in I} \mu(\widehat{s})$ if and only if $s \in I$ but $s' \notin I$. These are the only events in which DM expects to choose $c_{-s'}^{+s}$ from $\{c, c_{-s'}^{+s}\}$. We have

$$V\left(\{c, c_{-s'}^{+s}\}\right) = \begin{cases} \mu(s) + (1 - \mu(s)) c & \{s, s'\} \nsubseteq I \text{ for all } I \in \mathcal{P} \\ c & \text{otherwise} \end{cases},$$

and $\{c\} \succeq \{c_{-s'}^{+s}\}$ if and only if $\{s, s'\} \subseteq I$ for some $I \in \mathcal{P}$.

By the first observation, Definition 7 (1) is equivalent to the condition $\sigma(\mu^1) = \sigma(\mu^2)$. By the second observation, Definition 7 (2) is equivalent to the condition that, for all s and s', and for any c such that $\{c\} \succeq_i \{c_{-s'}^{+s}\}$ for $i = 1, 2$, $\{s, s'\} \subseteq I$ for some $I \in \mathcal{P}^1$ implies $\{s, s'\} \subseteq I$ for some $I \in \mathcal{P}^2$. In other words, $\mathcal{P}^1$ is finer then $\mathcal{P}^2$.

(ii) In the context of Theorem 3, a mean-preserving spread of the information structures is equivalent to having a finer partition with the same μ. The result is thus an immediate corollary of Theorem 2.

B.5. Proof of Theorem 5

We start with the following claim that guarantees the existence of a utility representation for $\succeq^*$.

Claim 8. *There exists $U : \mathcal{K}(\mathcal{F}) \times [0,1] \to \mathbb{R}$ that represents $\succeq^*$.*

Proof. It is enough to show that there exists a countable $\succeq^*$−dense subset of $\mathcal{K}(\mathcal{F}) \times [0,1]$. Note that a countable $\succeq_t^*$−dense subset is given by $\{\{c\} \, | c \in \mathbb{Q} \cap [0,1]\}$. We claim that for some $\widehat{t} \in [0,1]$, $\Psi := \{\{c\} \, | c \in \mathbb{Q} \cap [0,1]\} \times \{\widehat{t}\}$ is a countable $\succeq^*$−dense subset in $\mathcal{K}(\mathcal{F}) \times [0,1]$. First, Ψ is obviously countable. Now we need to show that if $(F, t) \succ^* (G, t')$ then we can find $(\{c\}, \widehat{t}) \in \Psi$ such that $(F, t) \succ^* (\{c\}, \widehat{t}) \succ^* (G, t')$. By Axiom 6, $\{1\} \succeq_t^* F \succeq_t^* \{0\}$, and since $\succeq_t^*$ on $\{c\}$ is strictly monotone, there exists $\{c_F\} \sim_t^* F$. Using Axiom 9, $(F, t) \sim^* (\{c_F\}, t) \sim^* (\{c_F\}, \widehat{t}) \succ^* (\{c_G\}, \widehat{t}) \sim^* (\{c_G\}, t') \sim^* (G, t')$, which means that $c_F > c_G$ and that there is $c \in \mathbb{Q} \cap [0,1]$ with $c_F > c > c_G$, such that $(F, t) \succ^* (\{c\}, \widehat{t}) \succ^* (G, t')$. $\blacksquare$

Let U represent $\succeq^*$. For each t, denote by U_t the restriction of U to $(\cdot, t)$. Since both U_t and V_t represents $\succeq_t^*$, there exists a strictly monotone function Γ_t, such that $V_t = \Gamma_t \circ U_t$. But since $V_t(\{c\}) = V_{t'}(\{c\})$ for all $\{c\}$, for any t' we also have $V_{t'} = \Gamma_t \circ U_{t'}$, hence Γ_t is independent of t. Without loss of generality, we can take Γ to be the identity function, which means that the function V^* satisfying $V^*(F, t) = V_t(F)$ represents $\succeq^*$. By Axiom 9, $V^*(\{f\}, t) = V^*(\{f\}, t')$ for all $t, t' \in [0, 1]$, and by Axiom 8, $V^*(F, t) \geq V^*(F, t')$ for all $F \in \mathcal{K}(\mathcal{F})$ and $t \geq t'$.

Claim 9. *If $t \geq t'$ then $\succeq_t^*$ has more preference for flexibility than $\succeq_{t'}^*$.*

Proof. If $V^*\left(\{f\}, t\right) \geq V^*\left(F, t\right)$ then $V^*\left(\{f\}, t'\right) = V^*\left(\{f\}, t\right) \geq V^*\left(F, t'\right)$. ∎

Finally, by Theorem 2, $\succeq_t^*$ has more preference for flexibility than $\succeq_{t'}^*$ if and only if p_t is a mean preserving spread of $p_{t'}$. This concludes the proof of Theorem 5.

References

[1] Anscombe, F.J. and R. J. Aumann (1963), "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34, 199-205.

[2] Aumann, R. J. (1976), "Agreeing to Disagree," *The Annals of Statistics*, Vol. 4, No. 6, 1236-1239.

[3] Blackwell, D. (1951), "The Comparison of Experiments," in: J. Neyman, Editor, Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, Univ. of California Press, Berkeley, 93–102.

[4] Blackwell, D. (1953), "Equivalent Comparison of Experiments," *Ann. Math. Statist.*, 24, 265–272.

[5] Dekel, E., B. Lipman, and A. Rustichini (2001), "Representing Preferences with a Unique Subjective State Space," *Econometrica*, 69, 891–934.

[6] Dekel, E., B. Lipman, A. Rustichini, and T. Sarver (2007), "Representing Preferences with a Unique Subjective State Space: Corrigendum," *Econometrica*, 75, 591–600. (See also supplemental appendix.)

[7] De Olivera, H. (2012) "Axiomatic Foundations for Rational Inattention." mimeo

[8] Dillenberger, D., and P. Sadowski (2012a), "Generalized Partition and Subjective Filtration," mimeo.

[9] Dillenberger, D., and P. Sadowski (2012b), "Subjective Learning," PIER Working Paper number 12-007.

[10] Ergin, H., and T. Sarver (2010), "A Unique Costly Contemplation Representation," *Econometrica*, 78, 1285-1339.

[11] Gollier, C., (2001), "The Economics of Risk and Time," The MIT Press, 0-262-07215-7.

[12] Harsanyi, John. C., (1967), "Games with Incomplete Information Played by "Bayesian" Players, I-III," *Management Science*, Vol. 14, No. 3, 159-182.

[13] Higashi, Y., K. Hyogo, and N. Takeoka (2009), "Subjective Random Discounting and Intertemporal Choice," *Journal of Economic Theory*, 144 (3), 1015–1053.

[14] Hyogo, K. (2007), "A Subjective Model of Experimentation," *Journal of Economic Theory*, 133, 316–330.

[15] Kihlstrom, R. E., (1984), "A Bayesian Exposition of Blackwell's Theorem on the Comparison of Experiments," in Bayesian Models of Economic Theory (M. Boyer and R. E. Kihlstrom, Eds.), pp. 13-32, Elsevier, Amsterdam.

[16] Kreps, D. (1979), "A Representation Theorem for "Preference for Flexibility," *Econometrica*, 47, 565–577.

[17] Krishna, R. V., and P. Sadowski (2012), "Dynamic Preference for Flexibility," mimeo.

[18] Lleras, J.S. (2011), "Expected Knowledge," unpublished manuscript.

[19] Mihmn M., and M.K. Ozbek, (2012) "Decision Making with Rational Inattention," mimeo.

[20] Savage, J. (1954), "The Foundations of Statistics." New York: John Wiley & Sons.

[21] Takeoka, N. (2005), "Subjective Prior over Subjective States, Stochastic Choice, and Updating," unpublished manuscript.

[22] Takeoka, N. (2007), "Subjective Probability over a Subjective Decision Tree." *Journal of Economic Theory*, 136, 536–571.