Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper 12-036

"Generalized Partition and Subjective Filtration"

by

David Dillenberger and Philipp Sadowski

http://ssrn.com/abstract=2146710

Generalized partition and subjective filtration*

David Dillenberger† Philipp Sadowski‡

September 2012

Abstract

We study an individual who faces a dynamic decision problem in which the process of information arrival is unobserved by the analyst, and hence should be identified from observed choice data. An information structure is objectively describable if signals correspond to events of the objective state space. We derive a representation of preferences over menus of acts that captures the behavior of a Bayesian decision maker who expects to receive such signals. The class of information structures that can support such a representation generalizes the notion of a partition of the state space. The representation allows us to compare individuals in terms of the preciseness of their information structures without requiring that they share the same prior beliefs. We apply the model to study an individual who anticipates gradual resolution of uncertainty over time. Both the filtration (the timing of information arrival with the sequence of partitions it induces) and prior beliefs are uniquely identified.

Key words: Resolution of uncertainty, valuing binary bets more, generalized partition, subjective filtration.

1. Introduction

1.1. Motivation

A standard dynamic decision problem involves specifying the set of possible states of nature, the set of available actions, and the information structure, which is the set of possible signals about the states that are expected to arrive over time and the probability of each signal given a state. The idea is that some signals are expected to be observed by an individual before he takes his final decision on the exposure to risk or uncertainty. For example, one way to analyze a career selection is to assume that potential workers who attend specialized

*First version September 2011. The results in this paper previously appeared in Dillenberger, D., and P. Sadowski (2012). We thank David Ahn, Brendan Daley, Haluk Ergin, Itzhak Gilboa, Faruk Gul, Peter Landry, Wolfgang Pesendorfer, Todd Sarver, Andrei Savochkin, and Roee Teper for useful advice.

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

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

schools or training programs, may gradually learn about their skills and abilities that affect their future occupational choice. 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. In this case, we would like to identify the subjective information structure (which contains signals the analyst is unaware of) from observed choice behavior and tie signals to observable components, that is, describe them as events of the objective state space.

As we explain in more detail below, we take as primitive a preference relation over sets (or menus) of acts defined over an objective state space. The interpretation is that the decision maker (henceforth DM) initially chooses among menus and subsequently chooses an act from the menu. We derive a *generalized-partition representation*, where the set of possible signals in the underlying information structure corresponds to subsets of the objective state space. The representation can be interpreted as follows: the DM behaves as if he has beliefs about which event he might know at the time he chooses from the menu. For any event, he 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 from the menu the act that maximizes the corresponding expected utility.

The model can accommodate a variety of information structures that capture interesting types of learning processes. As the name suggests, the notion of generalized partition extends the notion of a set partition, according to which the DM learns which cell of a (subjective) partition contains the true state. In the case of a set partition, signals are deterministic; that is, for each state there is only one possible event that contains it.[1] Another example of a generalized partition is a random partition, where one of multiple partitions is randomly drawn and then an event in it is reported. A situation that may give rise to a random partition is an experiment with uncertainty about its precision. A sequential elimination of candidates, say during a recruiting process, may also lead to learning via a generalized partition; if k candidates out of n are to be eliminated in the first stage, then the resulting collection of events the DM might possibly learn is the set of all $(n-k)$-tuples. We characterize the types of learning that can be accommodated by a generalized partition.

A different situation where signals may not be deterministic arises if the DM is unsure about the exact time at which he will have to choose an act from the menu. In that case, distinct events that contain the same state may simply correspond to the information the DM expects to have at different points in time. Reinterpreting our domain such that the opportunity to choose from a menu arrives randomly over time, we derive a subjective filtration representation that captures a DM who anticipates gradual resolution of uncertainty

[1] Partitional learning is analyzed in Dillenberger, Lleras, Sadowski, and Takeoka (2012).

over time. According to this representation, the DM behaves as if he holds prior beliefs over the state space and has in mind a filtration indexed by continuous time (that is, a sequence of increasingly finer partitions of the state space). Using Bayes' law, the filtration together with the beliefs generate a subjective temporal lottery. Both the filtration (the timing of information arrival with the sequence of partitions it induces) and prior beliefs are uniquely identified.

The description of signals as events of the state space facilitates the behavioral comparison of two individuals in terms of the information they expect to receive, independently of their prior beliefs. Any generalized-partition representation induces a natural comparative measure of "expecting to learn more". For example, in the case of the subjective-filtration representation, one individual expects to learn earlier than another if his filtration is finer at any point in time. We identify the behavioral implications of these measures in terms of individuals' desire to retain future betting options.

1.2. Formal preview of results

Let S be a finite state space. An act is a mapping $f : S \to [0,1]$, where $[0,1]$ is interpreted as a utility space. Preferences are defined over sets (or menus) of acts. The interpretation is that the 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 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. The preference relation over menus of acts is thus the only primitive of the model, leaving the uncertainty that the DM faces, as well as his ultimate choice of an act, unmodeled.

Our starting point is a general representation of preferences over menus of acts, which was first derived in Dillenberger, Lleras, Sadowski, and Takeoka (2012, henceforth DLST); DLST show that a familiar set of axioms is equivalent to a subjective-learning representation, according to which the value of a menu of acts F is given by

$$V(F) = \int_{\Delta(S)} \max_{f \in F} \left(\textstyle\sum_{s \in S} f(s)\, \pi(s) \right) dp(\pi), \tag{1}$$

where $p(\cdot)$ is a unique probability measure on $\Delta(S)$, the space of all probability measures on S. In representation (1) signals are identified with posterior distributions over S.

In this paper we study a specialized model in which signals are subsets of the state space, that is, elements of 2^S. We maintain the basic axioms of DLST and impose two additional axioms, Finiteness and Context independence. *Finiteness* implies that the probability measure

p in (1) has a finite support. (Finiteness is obviously necessary since 2^S is finite.) *Context independence* captures an idea that resembles Savage's (1954) sure-thing principle: if $f \neq g$ only on event I, and if g is unconditionally preferred to f, then the DM would also prefer g to f contingent upon learning I. The implication of this property in a dynamic decision problem is that if the DM prefers the singleton menu $\{g\}$ to $\{f\}$, then the DM would prefer to replace f with g on any menu $F \ni f$, from which he will choose f only if he learns I. We identify through preferences a special subset of menus, which we term saturated (Definition 4). The properties of a saturated menu $F \ni f$ are consistent with the interpretation that the DM anticipates choosing f from F only contingent on the event $\{s \in S \,|\, f(s) > 0\}$. Context independence requires that if $g(s) > 0 \Leftrightarrow f(s) > 0$ and $\{g\}$ is preferred over $\{f\}$, then the DM would prefer to replace f with g on any saturated menu $F \ni f$.

With these additional axioms, Theorem 1 derives a *generalized-partition representation* in which the value of a menu F is given by

$$V(F) = \textstyle\sum_{I \in 2^{\sigma(\mu)}} \max_{f \in F} \left[\textstyle\sum_{s \in I} f(s)\, \mu(s)\right] \rho(I),$$

where μ is a probability measure on S with support $\sigma(\mu)$, and $\rho : 2^{\sigma(\mu)} \to [0,1]$ is such that for any $s \in \sigma(\mu)$, ρ_s defined by $\rho_s(I) = \begin{cases} \rho(I) & \text{if } s \in I \\ 0 & \text{if } s \notin I \end{cases}$ satisfies $\sum_{I \subseteq \sigma(\mu)} \rho_s(I) = 1$. The pair (μ, ρ) is unique. We call the function ρ, which specifies the subjective information structure, a *generalized partition.* The probability of being in event I when the state of the world is s, $\rho_s(I)$, is the same for all states $s \in I$. This suggests that the DM is Bayesian and can only infer which states were excluded. In other words, the relative probability of any two states within an event is not updated.

The support of a generalized partition ρ is the set of possible signals and can be interpreted as a *type* of learning. We characterize all collections of events $\Psi \subseteq 2^S$ for which there is a generalized partition ρ with support Ψ. Theorem 2 shows that a necessary and sufficient condition is that Ψ be a uniform cover; we say that $\Psi \subseteq 2^S$ is a uniform cover of a set $S' \subseteq S$ if there exists $k \geq 1$ and a function $\beta : \Psi \to \mathbb{Z}_+$ such that for all $s \in S'$, $\sum_{I \in \Psi | s \in I} \beta(I) = k$. In this case we say that S' is covered k times by Ψ. Note that a set partition is implied if $k = 1$. The notion of uniform cover is closely related to the notion of a balanced collection of weights, as introduced by Shapley (1967) in the context of cooperative games.

As an example of a particular type of learning that can be accommodated in our framework, we show that the domain of menus of acts can capture a DM who expects uncertainty to resolve gradually over time. To this end, we reinterpret menus as choice situations in which the opportunity to choose from the menu arrives randomly. In this context, we interpret two events in the support of ρ in Theorem 1, both containing the sate s, as relevant for the DM at different points in time. If information becomes more precise over time, then the two events should be ordered by set inclusion. We use the notion of saturated menus to impose an additional axiom, *Sequential learning*, which formalizes this requirement. Theorem 3 provides a (subjective filtration) representation in which the value of a menu F is given by

$$V(F) = \int_{[0,1]} \left(\sum_{I \in \mathcal{P}_t} \max_{f \in F} \left[\sum_{s \in S} f(s) \mu(s) \right] \right) dt,$$

where μ is a probability measure on S and $\{\mathcal{P}_t\}$ is a filtration indexed by $t \in [0,1]$. The pair $(\mu, \{\mathcal{P}_t\})$ is unique.

Lastly, we use our representation results to compare the behavior of two individuals who expect to receive different information. These individuals differ in the value they derive from the availability of binary bets as intermediate actions. Suppose both DM1 and DM2 are sure to receive a certain payoff independently of the true state of the world. Roughly speaking, DM1 values binary bets more than DM2 if for any two states s and s', whenever DM1 prefers receiving additional payoffs in state s over having the option to bet on s versus s' (in the form of an act that pays well on s and nothing on s'), so does DM2. Theorem 4 states that in the context of Theorem 1 , DM1 *values binary bets more* than DM2 if and only if he expects to receive more information than DM2, in the sense that given the true state of the world, he is more likely to be able to rule out any other state (i.e. to learn an event, which contains the true state but not the other state.) In the context of Theorem 3, Theorem 5 first shows that DM1 values binary bets more than DM2 if and only if $\{\mathcal{P}_t^1\}$ is finer than $\{\mathcal{P}_t^2\}$ (i.e., for any t, all events in $\mathcal{P}_t^2$ are measurable in $\mathcal{P}_t^1$). Furthermore, if also $\mu^1 = \mu^2$, then $\{\mathcal{P}_t^1\}$ is finer than $\{\mathcal{P}_t^2\}$ if and only if DM1 has more preference for flexibility than DM2, in the sense that whenever DM1 prefers to commit to a particular action rather than to maintain multiple options, so does DM2.

1.3. Related literature

As mentioned above, the subjective-learning representation in DLST (2012) is the starting point for our analysis. Their partitional-learning representation is a special case of the model outlined in Section 2. It can also be viewed as a special case of the model in Section 3, where the DM does not expect to learn gradually over time, that is, he forms his final beliefs at time zero, right after he chose a menu. DLST further suggest how to study subjective temporal resolution of uncertainty by explicitly including the timing of the choice of an act in the domain. Takeoka (2007) studies subjective temporal resolution of uncertainty by analyzing choice between what one might term "compound menus" (menus over menus etc.). In Section 3.2 we compare those two approaches to ours.

More generally, our work is part of the literature on preference over sets of alternatives which assumes set monotonicity, that is, larger sets are weakly better than smaller ones (see, for example, Kreps (1979) and Dekel, Lipman, and Rustichini (2001)). Most papers in this literature study uncertainty over future tastes, and not over beliefs on an objective state space. Hyogo (2007), Ergin and Sarver (2010), De Olivera (2012), and Mihm and Ozbek (2012) use preferences over sets of alternatives to study models in which the DM can take costly actions that affect the process of information acquisition (either about his future tastes or future beliefs). In our paper the preciseness of information is not a choice variable, but a preference parameter that can be identified from choice data.

2. Subjective learning with unambiguously describable signals

Let $S = \{s_1, ..., s_k\}$ be a finite state space. An act is a mapping $f : S \to [0,1]$. Let $\mathcal{F}$ be the set of all acts. Let $\mathcal{K}(\mathcal{F})$ be the set of all non-empty compact subsets of $\mathcal{F}$. Capital letters denote sets, or menus, and small letters denote acts. For example, a typical menu is $F = \{f, g, h, ...\} \in \mathcal{K}(\mathcal{F})$. 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.[2] 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.

DLST (Theorem 1) derive the following representation of $\succeq$.

Definition 1. *A subjective-learning representation is a function $V : \mathcal{K}(\mathcal{F}) \to \mathbb{R}$, such that*

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

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

The axioms that are equivalent to the existence of a subjective-learning representation are familiar from the literature on preferences over menus of lotteries – *Ranking, vNM Continuity,*

[2] 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 from deriving the utility function (which is a standard exercise) by looking directly at utility acts instead of the corresponding Anscombe-Aumann acts.

Nontriviality, and *Independence* – adapted to the domain $\mathcal{K}(\mathcal{F})$, in addition to *Dominance*, which implies monotonicity in payoffs, and *Set Monotonicity*, which captures preference for flexibility.

A subjective-learning representation does not allow us to identify information independently of the induced changes in beliefs. The reason is that the signals are not objectively describable, that is, they are not stated in terms of the objective state space S. We thus study a model in which signals in a subjective-learning representation are restricted to correspond to events, that is, to subsets of S. The DM's beliefs can then be understood as uncertainty about the event he will know at the time of choosing from the menu. Throughout this section we assume that the relation $\succeq$ admits a subjective-learning representation. Section 2.1 develops a language that allows us to formulate a behavioral axiom, which implies that the DM cannot draw any inferences from learning an event besides knowing that states outside that event were not realized. Section 2.2 derives the most general representation in which signals correspond to events and the relative probability of any two states is the same across all events that contain them. A set of possible signals can be interpreted as a particular type of learning. Section 2.3 characterizes the types of learning that can be accommodated.

Since there are only finitely many distinct subsets of S, the support of the function p, denoted $\sigma(p)$, in a subjective-learning representation must be finite. This restriction is captured by the following axiom, which we also maintain throughout this section:

Axiom 1 (Finiteness). *For all $F \in \mathcal{K}(\mathcal{F})$, there is a finite set $G \subseteq F$ with $G \sim F$.*

The intuition for why Axiom 1 indeed implies that $\sigma(p)$ is finite is clear: if for any F there is a finite subset G of F that is as good as F itself, then only a finite set of posterior beliefs can be relevant. The formal statement of this result is provided by Riella (2011, Theorem 2), who establishes that Axiom 1 is the appropriate relaxation of the finiteness assumption in Dekel, Lipman, and Rustichini (2009, Axiom 11) if preferences are, as in this paper, monotonic with respect to set inclusion.

2.1. Axiom Context independence

The axiom we propose in this section captures an idea that resembles Savage's (1954) sure-thing principle: if $f \neq g$ only on event I, and if g is unconditionally preferred to f (that is, $\{g\} \succeq \{f\}$), then the DM would also prefer g to f contingent upon learning I. Since learning is subjective, stating the axiom requires us to first identify how the ranking of acts contingent on learning an event affects choice over menus. To this end, we now introduce the notion of saturated menus.

Definition 2. *Given $f \in \mathcal{F}$, let f_s^x be the act*

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

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

Definition 3. *A menu $F \in \mathcal{K}(\mathcal{F})$ is fat free if for all $f \in F$ and for all $s \in \sigma(f)$, $F \succ (F \backslash \{f\}) \cup \{f_s^0\}$.*

If a menu F is fat free, then for any act $f \in F$ and any state $s \in \sigma(f)$, eliminating s from $\sigma(f)$ reduces the value of the menu.[3] In particular, removing an act f from the fat-free menu F must make the menu strictly worse.

Definition 4. *A menu $F \in \mathcal{K}(\mathcal{F})$ is saturated if it is fat free and satisfies*

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

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

Definition 4 says that if F is a saturated menu, then (i) if an act $f \in F$ does not yield any payoff in some state, then the DM's preferences are insensitive to slightly improving f in that state; and, (ii) adding a collection of acts to a saturated menu implies that there is at least one act in the new menu that is not valued by the DM. In particular, the extended menu is no longer fat-free.

To better understand the notions of fat-free and saturated menus, consider the following example.

Example 1. *Suppose that there are two states $S = \{s_1, s_2\}$. If the act f yields positive payoffs in both states but only one of them is non-null,[4] then $\{f\}$ is not fat-free. If both states are non-null and f does not yield positive payoffs on one of them, then the set $\{f\}$ is not saturated according to Definition 4 (i). If the two states are non-null and f yields positive payoffs in both, then $\{f\}$ is fat-free, but it is not necessarily saturated. For example, if the DM expects to learn the true state for sure, that is, $\sigma^1(p) = \{(1,0), (0,1)\}$, then for $\varepsilon > 0$ and $g = (f(s_1) + \varepsilon, 0)$, both $\{f, g\} \succ \{f\}$ and $\{f, g\} \succ \{g\}$, which means that $\{f\}$ is not saturated according to Definition 4 (ii). On the other hand, if the DM expects to learn*

---

[3] Our notion resembles the notion of "fat-free acts" suggested by Lehrer (2012). An act f is fat-free if when an outcome assigned by f to a state is replaced by a worse one, the resulting act is strictly inferior to f. In our setting, a finite fat-free set contains acts, for all of which reducing an outcome in any state in the support results in an inferior set.

[4] In the context of a subjective-learning representation, a state s is non-null if $\int_{\Delta(S)} \pi(s)\, dp(\pi) > 0$.

nothing, that is $|\sigma^1(p)| = 1$*, then for all* g*, either* $\{f, g\} \sim \{f\}$ *or* $\{f, g\} \sim \{g\}$*, which means that* $\{f\}$ *is saturated.*

Claim 1. *A saturated menu* F*, with* $f(s) < 1$ *for all* $f \in F$ *and all* $s \in S$*, always exists. Furthermore, if* F *is saturated, then* F *is finite.*

Proof. See Appendix 6.1 ■

The following two claims illustrate properties of saturated menus in the context of a subjective-learning representation. In all that follows, we only consider saturated menus that consist of acts f with $f(s) < 1$ for all $s \in S$. For ease of exposition, we refrain from always explicitly stating this assumption.

Claim 2. *If* F *is saturated, then there is a one-to-one correspondence between* F *and the set of posterior beliefs.*

Proof. See Appendix 6.2 ■

Claim 2 connects the definition of a saturated menu with the idea that the DM might be required to make a decision when his state of knowledge is any one of the posterior beliefs from a subjective-learning representation. Claim 2 says that any act in a saturated menu is expected to be chosen under exactly one such belief.

The next claim demonstrates that the support of any act in a saturated menu coincides with that of the belief under which the act is chosen. For any act f in a given saturated menu F, let $\pi_f \in \sigma(p)$ be the belief such that $f = \arg\max_{f' \in F} \sum_{s \in S} f'(s) \pi_f(s)$. By Claim 2, π_f exists and is unique.

Claim 3. *If* F *is saturated and* $f \in F$ *then* $\sigma(f) = \sigma(\pi_f)$*.*

Proof. If $f(s) > 0$ and $\pi_f(s) = 0$, then $F \sim (F \backslash \{f\}) \cup \{f_s^0\}$, which is a contradiction to F being fat-free (and, therefore, saturated.) If $f(s) = 0$ and $\pi_f(s) > 0$, then for any $\varepsilon > 0$, $F \prec F \cup \{f_s^\varepsilon\}$, which is a contradiction to F being saturated. ■

We are now ready to state the central axiom of this section.

Axiom 2 (Context independence). *Suppose* F *is saturated and* $f \in F$*. Then for all* g *with* $\sigma(g) = \sigma(f)$*,*

$$\{g\} \succeq \{f\} \text{ implies } (F \backslash \{f\}) \cup \{g\} \succeq F.$$

Suppose the DM prefers committing to g over committing to f, where both g and f pay strictly positive amounts only on the event $\sigma(f)$. The axiom then requires that the DM would prefer to replace f with g on any saturated menu that contains f. To motivate this axiom, note that Claim 3 suggests that from a saturated menu $F \ni f$, DM plans to choose f if and only if he learns $\sigma(f)$. We would like to assume that $\{g\} \succeq \{f\}$ and $\sigma(f) = \sigma(g)$ imply that g is preferred to f contingent on $\sigma(f)$. Hence, $(F \setminus \{f\}) \cup \{g\} \succeq F$ should hold. This is Axiom 2.

2.2. Generalized-partition representation

Definition 5. *A function $\rho : 2^{S'} \to [0,1]$ is a generalized partition of $S' \subseteq S$ if for any $s \in S'$, ρ_s defined by $\rho_s(I) = \begin{cases} \rho(I) & \text{if } s \in I \\ 0 & \text{if } s \notin I \end{cases}$ satisfies $\sum_{I \subseteq S'} \rho_s(I) = 1$.*

We interpret $\rho_s(I)$ as the probability of signal I contingent on the state being s. The special case of a set partition corresponds to ρ taking only two values, zero and one. In that case, for every $s \in S'$ there exists a unique $I_s \in 2^{S'}$ with $s \in I_s$ and $\rho_s(I_s) = 1$. Furthermore, $s' \in I_s$ implies that $I_s = I_{s'}$, that is, $\rho_{s'}(I_s) = 1$ for all $s' \in I_s$.[5]

Definition 6. *The pair (μ, ρ) is a generalized-partition representation if (i) $\mu : S \to [0,1]$ is a probability measure; (ii) $\rho : 2^{\sigma(\mu)} \to [0,1]$ is a generalized partition of $\sigma(\mu)$; and (iii)*

$$V(F) = \sum_{I \in 2^{\sigma(\mu)}} \max_{f \in F} \left[\sum_{s \in I} f(s)\,\mu(s)\right] \rho(I)$$

represents $\succeq$.

The fact that $\rho_s(I)$ is independent of s (conditional on $s \in I$) reflects the idea that the DM cannot draw any inferences from learning an event other than that states outside that event were not realized. Indeed, Bayes' law implies that for any $s, s' \in I$,

$$\frac{\Pr(s\,|I)}{\Pr(s'\,|I)} = \frac{\rho_s(I)\,\mu(s)/\mu(I)}{\rho_{s'}(I)\,\mu(s')/\mu(I)} = \frac{\mu(s)}{\mu(s')} \tag{2}$$

independent of I. In that sense, the signals that support a generalized partition can be objectively described as events in $\sigma(\mu)$. It is worth noting that the notion of generalized partition is meaningful also in the context of objective learning, that is, when the function ρ is exogenously given.

Assume that the relation $\succeq$ admits a subjective-learning representation. We now show that further imposing Axiom 1 and Axiom 2 is both necessary and sufficient for having a generalized-partition representation.

[5] If ρ is partitional, then it is uniquely identified via its support, $\sigma(\rho)$. Throughout the paper, we use ρ and $\sigma(\rho)$ interchangeably when referring to a set partition.

Theorem 1. *Suppose that the relation $\succeq$ admits a subjective-learning representation (as in (1)). Then $\succeq$ satisfies Axioms 1 and 2 if and only if it has a generalized-partition representation, (μ, ρ). Furthermore, the pair (μ, ρ) is unique.*

Proof. See Appendix 6.3. ■

2.3. A characterization of generalized partitions

Definition 6 implies that (μ, ρ) is a generalized-partition representation if and only if ρ is a generalized partition of $\sigma(\mu)$. Equation (2) of the previous section observes that ρ is a generalized partition of $\sigma(\mu)$ if and only if the signals in its support can be objectively described as events in $\sigma(\mu)$. As we illustrate below, the set of events that are supported as possible signals can be interpreted as a particular type of learning. We now characterize the types of learning that can give rise to a generalized partition. Formally, we characterize the set

$$\left\{ \Psi \subseteq 2^{S'} \middle| \text{there is a generalized partition } \rho : 2^{S'} \to [0,1] \text{ with } \sigma(\rho) = \Psi \right\}.$$

Definition 7. *A set $S' \subseteq S$ is covered k times by a collection of events $\Psi \subseteq 2^S$ if there is a function $\beta : \Psi \to \mathbb{Z}_+$, such that for all $s \in S'$, $\sum_{I \in \Psi | s \in I} \beta(I) = k$.*

Definition 8. *A collection of events $\Psi \subseteq 2^S$ is a uniform cover of a set $S' \subseteq S$, if (i) $S' = \bigcup_{I \in \Psi} I$; and (ii) there exists $k \geq 1$, such that S' is covered k times by Ψ.*

Remark 1. *In the context of cooperative games, Shapley (1967) introduces the notion of a balanced collection of weights. Denote by $\mathcal{C}$ the set of all coalitions (subsets of the set N of players). The collection $(\gamma_L)_{L \in \mathcal{C}}$ of numbers in $[0,1]$ is a balanced collection of weights if for every player $i \in N$, the sum of γ_L over all the coalitions that contain i is 1. Suppose $\Psi \subseteq 2^S$ is a uniform cover of a set $S' \subseteq S$. Then there exists $k \geq 1$ such that for all $s \in S'$, $\sum_{I \in \Psi | s \in I} \frac{\beta(I)}{k} = 1$. In the terminology of Shapley, the collection $\left(\frac{\beta(I)}{k}\right)_{I \in \Psi}$ of numbers in $[0,1]$ is, thus, a balanced collection of weights.*

To better understand the notion of uniform cover, consider the following example. Suppose $S = \{s_1, s_2, s_3\}$. Any partition of S, for example $\{\{s_1\}, \{s_2, s_3\}\}$, is a uniform cover of S (with $k = 1$). A set that consists of multiple partitions, for example $\{\{s_1\}, \{s_2, s_3\}, \{s_1, s_2, s_3\}\}$, is a uniform cover of S (in this example with $k = 2$). The set $\Psi = \{\{s_2, s_3\}, \{s_1, s_2, s_3\}\}$ is *not* a uniform cover of S, because $\sum_{I | s_1 \in I} \beta(I) < \sum_{I | s_2 \in I} \beta(I)$ for any $\beta : \Psi \to \mathbb{Z}_+$. The set $\{\{s_2, s_3\}, \{s_1\}, \{s_2\}, \{s_3\}\}$, however, is a uniform cover of S with

$$\beta(I) = \begin{cases} 2 & \text{if } I = \{s_1\} \\ 1 & \text{otherwise} \end{cases}.$$

Lastly, the set $\{\{s_1, s_2\}, \{s_2, s_3\}, \{s_1, s_3\}\}$ is a uniform cover of S (with $k = 2$), even though it does not contain a partition.

An empirical situation that gives rise to a uniform cover consisting of two partitions is an experiment that reveals the state of the world if it succeeds, and is completely uninformative otherwise. For a concrete example that gives rise to a uniform cover that does not contain a partition, consider the sequential elimination of n candidates, say during a recruiting process. If k candidates are to be eliminated in the first stage, then the resulting uniform cover is the set of all $(n-k)$-tuples.

Theorem 2. *A collection of events Ψ is a uniform cover of $S' \subseteq S$ if and only if there is a generalized partition $\rho : 2^{S'} \to [0,1]$ with $\sigma(\rho) = \Psi$.*

Proof. See Appendix 6.4 ■

Theorem 2 characterizes the types of learning that can be accommodated by a generalized-partition representation. To illustrate it, let us consider a specific example. An oil company is trying to learn whether there is oil in a particular location. Suppose the company can perform a test-drill to determine accurately whether there is oil, $s = 1$, or not, $s = 0$. In that case, the company learns the partition $\{\{0\}, \{1\}\}$, and $\rho(\{0\}) = \rho(\{1\}) = 1$ provides a generalized-partition representation given the firm's prior beliefs μ on $S = \{0, 1\}$.

Now suppose that there is a positive probability that the test may not be completed (for some exogenous reason, which is not indicative of whether there is oil or not). The company will either face the trivial partition $\{\{0, 1\}\}$, or the partition $\{\{0\}, \{1\}\}$, and hence $\Psi = \{\{0, 1\}, \{0\}, \{1\}\}$. Suppose the company believes that the experiment will succeed with probability q. Then $\rho(\{0, 1\}) = 1 - q$ and $\rho(\{0\}) = \rho(\{1\}) = q$ provides a generalized-partition representation given the company's prior beliefs μ on $S = \{0, 1\}$.

We can extend the previous example and suppose the company is trying to assess the size of an oil field by drilling in l proximate locations, which means that the state space is now $\{0, 1\}^l$. As before, any test may not be completed, independently of the other tests. This is an example of a situation where the state consists of l different attributes (i.e., the state space is a product space), and the DM may learn independently about any of them. Such learning about attributes also gives rise to a uniform cover that consists of multiple partitions and can be accommodated.

To find a generalized-partition representation based on (i) a uniform cover Ψ of a state space S, for which there is a collection Π of partitions whose union is Ψ; (ii) a probability distribution q on Π; and (iii) a measure μ on S, one can set $\rho(I) = \sum_{\mathcal{P} \in \Pi | I \in \mathcal{P}} q(\mathcal{P})$. We refer to the pair (q, Π) as a random partition.

Lastly, reconsider the example of sequential elimination of candidates outlined above. Suppose that there are three candidates and one of them will be eliminated in the first round. The state space is then $S = \{s_1, s_2, s_3\}$ and the uniform cover of events the DM might learn is given by $\Psi = \{\{s_1, s_2\}, \{s_2, s_3\}, \{s_1, s_3\}\}$. Suppose that, contingent on person i being the best candidate, the DM considers any order of elimination of the other candidates as equally likely. This corresponds to the generalized partition with $\rho(I) = 0.5$ for all $I \in \Psi$ and $\rho(I) = 0$ otherwise. The pair (μ, ρ) is then a generalized-partition representation for any prior beliefs μ with full support on S.

3. Subjective temporal resolution of uncertainty

Suppose that the DM anticipates uncertainty to resolve gradually over time. The pattern of resolution might be relevant if, for example, the time at which the DM has to choose an alternative from the menu is random and continuously distributed over some interval, say $[0, 1]$. An alternative interpretation is that at any given point in time $t \in [0, 1]$ the DM chooses one act from the menu. At time 1, the true state of the world becomes objectively known. The DM is then paid the convex combination of the payoffs specified by all acts on the menu, where the weight assigned to each act is simply the amount of time the DM held it. That is, the DM derives a utility flow from holding a particular act, where the state-dependent flow is determined ex-post, at the point when payments are made. In both cases, the information available to the DM at any point in time t might be relevant for his choice. This section is phrased in terms of random timing of second-stage choice.

In a context where the flow of information over time is objectively given, it is common to describe it as a filtered probability space, that is, a probability space with a filtration on its sigma algebra. We would like to replicate this description in the context of subjective learning. We interpret two events in the support of ρ in Theorem 1, both containing state s, as relevant for the DM at different points in time. If signals arrive sequentially, then information becomes more precise over time, that is, the two events should be ordered by set inclusion. Using the notion of saturated menus, we now impose an additional axiom on $\succeq$, which captures this restriction. The resulting representation can be interpreted as follows: the DM holds beliefs over the states of the world and has in mind a filtration indexed by continuous time. Using Bayes' law, the filtration and prior beliefs jointly generate a subjective temporal lottery. Our domain is rich enough to allow both the filtration, that is the timing of information arrival and the sequence of partitions induced by it, and the beliefs to be uniquely identified from choice behavior.

3.1. Subjective filtration

Definition 9. *An act f contains act g if $\sigma(g) \subsetneq \sigma(f)$.*

Definition 10. *Acts f and g do not intersect if $\sigma(g) \cap \sigma(f) = \emptyset$.*

Axiom 3 (Sequential learning). *If F is saturated and $f, g \in F$ then either f and g do not intersect or one contains the other.*

As we explained above, in order to interpret two distinct events that contain state s as being relevant for the DM at different points in time, they must be ordered by set inclusion. Using Claim 3, this is the content of Axiom 3.

Definition 11. *The pair $(\mu, \{\mathcal{P}_t\})$ is a subjective filtration representation if (i) μ is a probability measure on S; (ii) $\{\mathcal{P}_t\}$ is a filtration on $\sigma(\mu)$ indexed by $t \in [0,1]$;[6] and*

$$V(F) = \int_{[0,1]} \left\{ \sum_{I \in \mathcal{P}_t} \max_{f \in F} \left[\sum_{s \in I} f(s) \mu(s) \right] \right\} dt.$$

represents $\succeq$.

Note that there can only be a finite number of times at which the filtration $\{\mathcal{P}_t\}$ becomes strictly finer. The definition does not require $\mathcal{P}_0 = \{\sigma(\mu)\}$.

Theorem 3. *Suppose that the relation $\succeq$ admits a Generalized-partition representation (as in Definition 6). The relation $\succeq$ satisfies Axiom 3 if and only if it has a subjective filtration representation, $(\mu, \{\mathcal{P}_t\})$. Furthermore, the pair $(\mu, \{\mathcal{P}_t\})$ is unique.*

Proof. See Appendix 6.5 ■

If the DM faces an (exogenously given) random stopping time that is uniformly distributed on $[0,1]$,[7] then Theorem 3 implies that he behaves as if he holds prior beliefs μ and expects to learn over time as described by the filtration $\{\mathcal{P}_t\}$.

We now briefly sketch the proof of Theorem 3. Given a generalized-partition representation (μ, ρ) as in Theorem 1, Axiom 3 allows us to construct a random partition (q, Π) as defined at the end of Section 2.3, where the partitions in Π can be ordered by increasing

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

[7] It is straightforward to accommodate any other exogenous distribution of stopping times. An alternative interpretation is that the distribution of stopping times is not uniform because of an external constraint, but because the DM subscribes to the principle of insufficient reason, by which he assumes that all points in time are equally likely to be relevant for choice.

fineness. If the DM faces a random stopping time that is uniformly distributed on $[0,1]$, then we can interpret $q(\mathcal{P})$ as the time for which the DM expects partition $\mathcal{P} \in \Pi$ to be relevant. This interpretation is captured in the time dependency of $\{\mathcal{P}_t\}$. The construction of (q,Π) is recursive. First, for each state $s \in S$, we find the largest set in $\sigma(\rho)$ that includes s. The collection of those sets constitutes $\mathcal{P}_1$. The probability $q(\mathcal{P}_1)$ corresponds to the smallest weight any of those sets is assigned by ρ, that is, $q(\mathcal{P}_1) = \min_{I \in \mathcal{P}_1}(\rho(I))$. To begin the next step, we calculate adjusted weights, ρ_1, as follows: for any set $I \in \mathcal{P}_1$, let $\rho_1(I) = \rho(I) - q(\mathcal{P}_1)$. For any set $I \in \sigma(\rho) \backslash \mathcal{P}_1$, let $\rho_1(I) = \rho(I)$. The set $\sigma(\rho_1)$ then consists of all sets $I \in \mathcal{P}_1$ with $\rho(I) > q(\mathcal{P}_1)$ and all sets in $\sigma(\rho) \backslash \mathcal{P}_1$. Recursively, construct $\mathcal{P}_n$ according to ρ_{n-1}. By Theorem 1, $\sum_{I \in 2^S | s \in I} \rho(I) = 1$ for all $s \in \sigma(\mu)$, which guarantees that the inductive procedure is well defined. It must terminate in a finite number of steps due to the finiteness of 2^S.

Remark 2. *At the time of menu choice, the DM holds beliefs over all possible states of the world. If he expects additional information to arrive before time-zero (at which point his beliefs commence to be relevant for choice from the menu), then time-zero information is described by a non-trivial partition of $\sigma(\mu)$, that is, $\mathcal{P}_0 \neq \{\sigma(\mu)\}$. If one wants to assume that the (subjective) flow of information cannot start before time-zero, then the following additional axiom is required:*

Axiom 4 (Initial node). *If F is saturated, then there exists $f \in F$ such that f contains g for all $g \in F$ with $g \neq f$.*

Under the assumptions of Theorem 3, if $\succeq$ also satisfies Axiom 4, then $\mathcal{P}_0 = \{\sigma(\mu)\}$. That is, the tree $(\mu, \{\mathcal{P}_t\})$ has a unique initial node (see Claim 10 in Appendix 6.5).

3.2. Reevaluation of our domain

In this section we compare our model of subjective filtration to two different approaches that have been suggested in the literature to study a DM who expects uncertainty to be gradually resolved over time.

Takeoka (2007) analyzes choice between what one might term "compound menus of acts" (menus over menus etc.), that is $\mathcal{K}(\mathcal{K}...\mathcal{K}(\mathcal{F}))$. The domain of compound menus provides a way to talk about compound uncertainty (without objective probabilities). It has the advantage that it can capture situations where the DM faces intertemporal trade-offs, for example if today's action may limit tomorrow's choices. However, while only the initial choice is modeled explicitly, the interpretation of choice on this domain now involves multiple stages, say 0, 1/2, and 1, at which the DM must make a decision. That is, the pattern of information arrival (or, at least, the collection of times at which an outside observer can detect changes in the DM's beliefs) is objectively given. In that sense, the domain only partially captures subjective temporal resolution of uncertainty. Furthermore, the domain of compound menus becomes increasingly complicated, as the resolution of uncertainty becomes finer.[8]

DLST (2012) extend the domain of menus of acts to one in which 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. More formally, they consider the domain $\mathcal{K}(\mathcal{F}) \times [0,1]$, where a typical element (F,t) represents a menu and a time by which an alternative from the menu must be chosen. DLST can accommodate intertemporal trade-offs, as perceived from the ex-ante point of view. For example, the DM can compare the alternatives (F,t) and (G,t'), where $t > t'$ and $F \subset G$; while anticipating late resolution of uncertainty provides an incentive to postpone his decision, the available menu at the later time will be worse. DLST provide a representation of preferences over $\mathcal{K}(\mathcal{F}) \times [0,1]$ that 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). While the dimensionality of $\mathcal{K}(\mathcal{F}) \times [0,1]$ is significantly lower than that of $\mathcal{K}(\mathcal{K}...\mathcal{K}(\mathcal{F}))$, it still requires the elicitation of the DM's ranking over menus of acts for every point in time.

In this paper we take a different approach to study subjective temporal resolution of uncertainty: we specify a single set of feasible intermediate actions, which is the relevant constraint on choice at *all* points in time. At the first stage, the DM chooses a menu of acts and only this choice is modeled explicitly. The innovation lies in our interpretation of choice from the menu. Whether we think of an exogenous distribution for the stopping time or of a model where the DM derives a utility flow, the information that the DM has at any point in time might be relevant for his ultimate choice from a menu. Our approach does not accommodate choice situations where the set of feasible actions may change over time. It also cannot capture the intertemporal trade-off as in DLST. Its main advantage, however, is that it allows us (as we argue in the text) to uniquely pin down the timing of information arrival in continuous time, the sequence of induced partitions, and the DM's prior beliefs from the familiar, much smaller, and analytically tractable domain of menus of acts.

4. Comparing valuations of binary bets

Under the assumptions of Theorem 1, we compare the behavior of two individuals in terms of the amount of information each expects to acquire.

[8] Note that the set of menus over acts is infinitely dimensional. Hence, even the three-stage model considers menus that are subsets of an infinite dimensional space.

Fix $k \in (0, 1-c)$ such that $\{c\} \succ_i \{f\}$ for $i = 1, 2$, where

$$f(\widehat{s}) = \begin{cases} c+k & \text{if } \widehat{s} = s \\ 0 & \text{if } \widehat{s} = s' \\ c & \text{otherwise} \end{cases}.$$

Let

$$f'(\widehat{s}) = \begin{cases} c+k' & \text{if } \widehat{s} = s \\ c & \text{otherwise} \end{cases}.$$

Relative to the certain payoff c, the act f is a bet with payoffs k in state s and $-c$ in state s'. The act f' yields extra payoff k' in state s.

Definition 12. *DM1 values binary bets more than DM2 if for all $s, s' \in S$ and $k' \in [0, k]$,*
(i) $\{f'\} \sim_1 \{c\} \Leftrightarrow \{f'\} \sim_2 \{c\}$; and
(ii) $\{f'\} \succeq_1 \{f, c\} \Rightarrow \{f'\} \succeq_2 \{f, c\}$.

Condition (i) says that the two DMs agree on whether or not payoffs in state s are valuable. Condition (ii) says that DM1 is willing to pay more in state s to have the bet f available.[9]

A natural measure of the amount of information that a DM expects to receive is how likely he expects to be able to distinguish any state s from any other state s' whenever s is indeed the true state. Observe that $\Pr(\{I \,|s \in I, \, s' \notin I\} \,|s) = \sum_{I|s \in I, s' \notin I} \rho(I)$.

Theorem 4. *If DM1 and DM2 have preferences that can be represented as in Theorem 1, then DM1 values binary bets more than DM2 if and only if $\sigma(\mu^1) = \sigma(\mu^2)$ and*

$$\sum_{I|s \in I, s' \notin I} \rho^1(I) \geq \sum_{I|s \in I, s' \notin I} \rho^2(I)$$

for all $s, s' \in \sigma(\mu^1)$.

Proof. See Appendix 6.6 ■

Consider the following definition of "more preference for flexibility" (see DLST and references therein), according to which 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] The notion of valuing binary bets more extends the notion of *valuing more binary bets*, which was introduced and characterized in DLST, in the context of partitional learning. While in both versions the two DMs should have the same set of states on which payoffs are valuable, the intensity of preferences does not play any role in DLST, while it does here.

Definition 13. *DM1 has more preference for flexibility than DM2 if for all $f \in \mathcal{F}$ and for all $G \in \mathcal{K}(\mathcal{F})$,*

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

DLST (Theorem 2) show that DM1 has more preference for flexibility than DM2 if and only if DM2's distribution of posterior beliefs (in a subjective-learning representation) is a mean-preserving spread of DM1's. Our notion of valuing binary bets more weakens the notion of more preference for flexibility.[10] Accordingly, Theorem 4 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.

Under the assumptions of Theorem 3, we can characterize the notion of preference for flexibility and the value of binary bets via the DM's subjective filtration.

Definition 14. *DM1 learns earlier than DM2 if $\{\mathcal{P}_t^1\}$ is weakly finer than $\{\mathcal{P}_t^2\}$.*

Theorem 5. *If DM1 and DM2 have preferences that can be represented as in Theorem 3, then:*

(i) DM1 values binary bets more than DM2 if and only if DM1 learns earlier than DM2;

(ii) DM1 has more preference for flexibility than DM2 if and only if DM1 learns earlier than DM2 and they have the same prior beliefs, $\mu^1 = \mu^2$.

Proof. See Appendix 6.7 ■

Theorem 5 shows that under the assumptions of Theorem 3, the characterization of more preference for flexibility differs from that of the weaker notion of valuing binary bets more solely by requiring that the prior beliefs are the same.

5. Concluding remarks

In this paper we show how to identify unobserved information instead of taking it as a primitive of a model. We provide minimal conditions under which a subjective information structure can be elicited from choice behavior and can be described solely in terms of the objective state space S. The set of possible signals in the corresponding generalized-partition representation corresponds to events, that is, subsets of S. The notion of generalized partition extends the notion of a set partition by dropping the requirement of deterministic signals.

[10] To see this, note that Condition (ii) in Definition 12 is implied by Definition 13 and Condition (i) in Definition 12 is implied by Claim 1 in DLST, which states that if DM1 has more preference for flexibility than DM2 then $\{f\} \succeq_1 \{g\}$ if and only if $\{f\} \succeq_2 \{g\}$.

We show that the model can accommodate a variety of information structures that induce interesting types of learning processes.

We now provide a simple example, the purpose of which is to demonstrate an empirically plausible behavioral pattern which our axioms accommodate, but is precluded under the stronger requirement of learning by a partition of S. Consider the state space $\{s_1, s_2\}$ and the menu $\{(1,0), (0,1), (1-\varepsilon, 1-\varepsilon)\}$, which contains the option to bet on either state, as well as an insurance option that pays reasonably well in both states. A DM who is uncertain about the information he will receive by the time he has to choose from the menu may strictly prefer this menu to any of its subsets (for ε small enough). For instance, an investor may value the option to make a risky investment in case he understands the economy well, but also values the option to make a safe investment in case a lot of uncertainty remains unresolved at the time of making the investment choice. Our axioms accommodate this ranking. In contrast, such a ranking of menus is ruled out if signals are deterministic. If the DM expects to learn the true state, then preference for flexibility stems *exclusively* from the DM's prior uncertainty about the true state and the insurance option is irrelevant, that is, $\{(1,0), (0,1), (1-\varepsilon, 1-\varepsilon)\} \sim \{(1,0), (0,1)\}$. And if the DM does not expect to learn the true state, then, for ε small enough, he anticipates choosing the insurance option with certainty, that is, $\{(1,0), (0,1), (1-\varepsilon, 1-\varepsilon)\} \sim \{(1-\varepsilon, 1-\varepsilon)\}$.

It is worth emphasizing the merit of taking the objective state space and the acts defined on it as primitives of the model. Unaware of the collection of random variables the DM might observe, an analyst may wish to follow Savage (1954) and postulate the existence of a grand state space that describes all *conceivable* sources of uncertainty. This expanded state-space surely captures all the uncertainty which is relevant for the DM. Identification of beliefs on a larger state space, however, generally requires a much larger collection of acts, which poses a serious conceptual problem, as in many applications the domain of choice (the available acts) is given. In that sense, acts – and the state space on which they are defined – should be part of the primitives of the model.[11] Instead of enlarging the state space, our approach identifies a behavioral criterion for checking whether a given state space (e.g. the one acts are naturally defined on in a particular application) is large enough: behavior satisfies Context independence (Axiom 2) if and only if the resolution of any subjective uncertainty corresponds to an event in the state space.

[11] Gilboa, Postlewaite, and Schmeidler (2009a, 2009b) point out the problems involved in using an analytical construction, according to which states are defined as functions from acts to outcomes, to generate a state space that captures all conceivable sources of uncertainty. First, since all possible acts on this new state space should be considered, the new state space must be extended yet again, and this iterative procedure does not converge. Second, the constructed state space may include events that are never revealed to the DM, and hence some of the comparisons between acts may not even be potentially observable. (A related discussion appears in Gilboa (2009, Section 11.1.)

6. Appendix

6.1. Proof of Claim 1

We will construct a menu that satisfies Definition 4 with $f(s) < 1$ for all $f \in F$ and all $s \in S$. Let $F_{\Delta(S)} := \{f \in \mathcal{F} : \|f\|_2 = 1\}$ be the positive segment of the $k-1$ dimensional unit sphere. There is an isomorphism between $\Delta(S)$ and $F_{\Delta(S)}$ with the mapping $\pi \to \underset{f \in F_{\Delta(S)}}{\arg\max} \left(\sum_{s \in S} f(s)\pi(s)\right)$. For $\mathcal{L} \subset \Delta(S)$ let $F_{\mathcal{L}} \subset F_{\Delta(S)}$ be the image of $\mathcal{L}$ under this mapping. Finiteness of $\sigma(p)$ implies that $F_{\sigma(p)}$ is finite. Let $f_{\sigma(p),\pi} := \underset{f \in F_{\sigma(p)}}{\arg\max} \left(\sum_{s \in S} f(s)\pi(s)\right)$ and (implicitly) define $\pi_{\sigma(p),f}$ by $f = \underset{f \in F_{\sigma(p)}}{\arg\max} \left(\sum_{s \in S} f(s)\pi_{\sigma(p),f}(s)\right)$. Because $F_{\Delta(S)}$ is the positive segment of a sphere, $\pi(s) > 0$ for $\pi \in \sigma(p)$ if and only if $f_{\sigma(p),\pi}(s) > 0$. This implies that $F_{\sigma(p)} \succ F_{\sigma(p)} \backslash \{f\} \cup \{f_s^0\}$ for all $f \in F_{\sigma(p)}$ and $s \in S$ with $f(s) > 0$. Hence, $F_{\sigma(p)}$ is fat-free (Definition 3). We claim that $F_{\sigma(p)}$ is a saturated menu. Consider condition (i) in Definition 4. If $f(s) = 0$, then $\pi_{\sigma(p),f}(s) = 0$. Hence, there exists $\bar{\varepsilon} > 0$ such that $F_{\sigma(p)} \sim F_{\sigma(p)} \cup \left\{f_s^{f(s)+\varepsilon}\right\}$ for all $\varepsilon < \bar{\varepsilon}$. Finally, consider condition (ii) in Definition 4. Let $G \nsubseteq F_{\sigma(p)}$. If $F_{\sigma(p)} \cup G \sim F_{\sigma(p)}$ then the condition is trivially satisfied. Suppose $F_{\sigma(p)} \cup G \succ F_{\sigma(p)}$. Then, there exist $\pi \in \sigma(p)$ and $g \in G$ with $\sum_{s \in S} g(s)\pi(s) > \sum_{s \in S} f_{\sigma(p),\pi}(s)\pi(s)$. Then $F_{\sigma(p)} \cup G \sim \left(F_{\sigma(p)} \cup G\right) \backslash \left\{f_{\sigma(p),\pi}\right\}$.

6.2. Proof of Claim 2

If F is saturated and $f \in F$, then there exists π such that $f = \arg\max\left(\sum_{s \in S} f(s)\pi(s)\right)$ (if not, then $F \sim F \backslash \{f\}$). We should show that if $f = \arg\max\left(\sum_{s \in S} f(s)\pi(s)\right)$, then for all $\pi' \neq \pi$, $f \notin \arg\max\left(\sum_{s \in S} f(s)\pi'(s)\right)$. Suppose to the contrary that there exist $\pi \neq \pi'$ such that $f = \arg\max\left(\sum_{s \in S} f(s)\pi(s)\right)$ and $f \in \arg\max\left(\sum_{s \in S} f(s)\pi'(s)\right)$. Then $f(s) > 0$ for all $s \in \sigma(\pi) \cup \sigma(\pi')$ by Definition 4 (i). We construct an act f', which does better than f with respect to belief π' and does not change the arg max with respect to any other belief in which f was not initially the best. Since $\pi \neq \pi'$, there exist two states, s and s', such that $\pi'(s) > \pi(s)$ and $\pi'(s') < \pi(s')$. Let

$$f'(\widehat{s}) = \begin{cases} f(\widehat{s}) \text{ if } \widehat{s} \notin \{s, s'\} \\ f(\widehat{s}) + \varepsilon \text{ if } \widehat{s} = s \\ f(\widehat{s}) - \delta \text{ if } \widehat{s} = s' \end{cases},$$

where $\varepsilon, \delta > 0$ are such that:
(1) $\varepsilon\pi'(s) - \delta\pi'(s') > 0$, and
(2) $\varepsilon\pi(s) - \delta\pi(s') < 0$.

The two conditions can be summarized as $\frac{\varepsilon}{\delta} \in \left(\frac{\pi'(s')}{\pi'(s)}, \frac{\pi(s')}{\pi(s)}\right) \subset (0, \infty)$. Note that one can make ε and δ sufficiently small (while maintaining their ratio fixed) so that, by continuity, f' does not change the arg max with respect to any other belief in which f was not initially the best. Hence $f' \notin F$ and $F \cup f' \succ (F \cup f') \setminus \{g\}$ for all $g \in F \cup f'$, which is a contradiction to F being saturated.

6.3. Proof of Theorem 1

Axiom 1 is obviously necessary for the representation. We thus only verify that the representation implies Axiom 2. Suppose then that F is saturated with $f \in F$, and let g satisfy $\sigma(g) = \sigma(f)$ and $\{g\} \succeq \{f\}$, which implies that

$$\begin{aligned} V(\{g\}) - V(\{f\}) &= \textstyle\sum_{I \in 2^{\sigma(\mu)}} \sum_{s \in I} [g(s) - f(s)]\, \mu(s)\, \rho(I) \\ &= \textstyle\sum_{s \in S} \sum_{I \in 2^{\sigma(\mu)} | s \in I} [g(s) - f(s)]\, \mu(s)\, \rho(I) \\ &= \textstyle\sum_{s \in S} [g(s) - f(s)]\, \mu(s) \sum_{I \in 2^{\sigma(\mu)} | s \in I} \rho(I) \\ &= \textstyle\sum_{s \in S} [g(s) - f(s)]\, \mu(s) \geq 0. \end{aligned} \quad (3)$$

Since F is saturated, Claim 2 and Claim 3 imply that there exists $I_f \in \sigma(\rho)$ such that

$$\begin{aligned} V(F) &= \left[\textstyle\sum_{s \in I_f} f(s)\, \mu(s)\right] \rho(I_f) + \textstyle\sum_{I \in 2^{\sigma(\mu)}/I_f} \max\limits_{f' \in F/\{f\}} \left[\sum_{s \in I} f(s)\, \mu(s)\right] \rho(I) \\ &\leq \left[\textstyle\sum_{s \in I_f} g(s)\, \mu(s)\right] \rho(I_f) + \textstyle\sum_{I \in 2^{\sigma(\mu)}/I_f} \max\limits_{f' \in F/\{f\}} \left[\sum_{s \in I} f(s)\, \mu(s)\right] \rho(I) \\ &\leq V((F \setminus \{f\}) \cup \{g\}), \end{aligned}$$

where the first inequality uses Equation (3) and the second inequality is because the addition of the act g might increase the value of the second component. Therefore, $(F \setminus \{f\}) \cup \{g\} \succeq F$.

The sufficiency part of Theorem 1 is proved using the following claims. The first establishes that a strict version of Axiom 2 also holds.

Claim 4. *Suppose F is saturated and $f \in F$. Then for all g with $\sigma(g) = \sigma(f)$,*

$$\{g\} \succ \{f\} \text{ implies } (F \setminus \{f\}) \cup \{g\} \succ F.$$

Proof. For $\varepsilon > 0$ small enough, let

$$h(s) = \begin{cases} f(s) + \varepsilon & \text{if } s \in \sigma(f) \\ 0 & \text{if } s \notin \sigma(f) \end{cases}.$$

Then $\{g\} \succ \{h\}$ and $\sigma(h) = \sigma(g)$. A subjective-learning representation implies that $F \cup \{h\} \succ F$. Let

$$F' := \left\{ \underset{f' \in F \cup \{h\}}{\arg\max} \left(\textstyle\sum_{s \in S} f'(s)\, \pi(s) \right) \middle| \, \pi \in \sigma(p) \right\}.$$

Then $F' \sim F \cup \{h\}$ and F' is saturated. By Axiom 2,

$$(F' \backslash \{h\}) \cup \{g\} \succeq F'.$$

Furthermore, $F' \backslash \{h\} \subseteq F \backslash \{f\}$ and, since any $\succeq$ that admits a subjective-learning representation is monotonic with respect to set inclusion, $(F \backslash \{f\}) \cup \{g\} \succeq (F' \backslash \{h\}) \cup \{g\}$. Collecting all the preference statements established above completes the proof:

$$(F \backslash \{f\}) \cup \{g\} \succeq (F' \backslash \{h\}) \cup \{g\} \succeq F' \sim F \cup \{h\} \succ F.$$

■

Claim 5. *If $\pi, \pi' \in \sigma(p)$ and $\pi \neq \pi'$ then $\sigma(\pi) \neq \sigma(\pi')$*

Proof. Suppose there are $\pi, \pi' \in \sigma(p)$, $\pi \neq \pi'$, but $\sigma(\pi) = \sigma(\pi')$. Let $F_{\sigma(p)}$ be the saturated menu constructed in the proof of Claim 1. Then there are $f, g \in F_{\sigma(p)}$ with $f \neq g$ but $\sigma(f) = \sigma(g)$. Without loss of generality, suppose that $\{g\} \succeq \{f\}$. For $\varepsilon > 0$ small enough, let

$$h(s) = \begin{cases} g(s) + \varepsilon & \text{if } s \in \sigma(f) \\ 0 & \text{if } s \notin \sigma(f) \end{cases}$$

and let

$$F := \left\{ \arg \max_{f \in F_{\sigma(p)} \cup \{h\}} \left(\textstyle\sum_{s \in S} f(s)\, \pi(s) \right) | \pi \in \sigma(p) \right\}.$$

F is a saturated menu with $F \sim F_{\sigma(p)} \cup \{h\}$. For $\varepsilon > 0$ small enough, $f, h \in F$. Furthermore, $\{h\} \succ \{g\} \succeq \{f\}$. Then, by Claim 4 $F \backslash \{f\} = (F \backslash \{f\}) \cup \{h\} \succ F$, which contradicts monotonicity with respect to set inclusion. ■

The measure p over $\Delta(S)$ in a subjective-learning representation is unique. Consequently the prior, $\mu(s) = \int_{\Delta(S)} \pi(s)\, dp$, is also unique. By Claim 5, we can index each element $\pi \in \sigma(p)$ by its support $\sigma(\pi) \in 2^S$ and denote a typical element by $\pi(\cdot | I)$, where $\pi(s | I) = 0$ if $s \notin I \in 2^S$. This allows us to replace the integral over $\Delta(S)$ with a summation over 2^S according to a unique measure $\widehat{p}$,

$$V(F) = \textstyle\sum_{I \in 2^S} \max_{f \in F} \left[\textstyle\sum_{s \in S} f(s)\, \pi(s | I) \right] \widehat{p}(I), \tag{4}$$

and to write $\mu(s) = \sum_{I|s\in I} \pi(s|I)\,\widehat{p}(I)$.

Claim 6. *For all* $s, s' \in I \in \sigma(\widehat{p})$,

$$\frac{\pi(s|I)}{\pi(s'|I)} = \frac{\mu(s)}{\mu(s')}.$$

Proof. Suppose to the contrary that there are $s, s' \in I \in \sigma(\widehat{p})$ such that

$$\frac{\pi(s|I)}{\pi(s'|I)} < \frac{\mu(s)}{\mu(s')}.$$

Given a saturated menu F, let $f_I := \arg\max_{f\in F} \sum_{\widehat{s}\in S} f(\widehat{s})\,\pi(\widehat{s}|I)$. By continuity, and since $f_I(s') > 0$, there exists an act h with

$$h(\widehat{s}) = \begin{cases} f_I(\widehat{s}) & \text{if } \widehat{s} \notin \{s, s'\} \\ f_I(\widehat{s}) + \varepsilon & \text{if } \widehat{s} = s \\ f_I(\widehat{s}) - \delta & \text{if } \widehat{s} = s' \end{cases},$$

where $\varepsilon, \delta > 0$ are such that:

(1) $\varepsilon\mu(s) - \delta\mu(s') > 0$, and

(2) $\varepsilon\pi(s|I) - \delta\pi(s'|I) < 0$.

Note that using Claim 2 and Claim 3 one can make ε and δ sufficiently small (while maintaining their ratio fixed), so that, by continuity and finiteness of $\sigma(\widehat{p})$, h does not change the arg max with respect to any other belief in $\sigma(\widehat{p})$. Then $\{h\} \succeq \{f_I\}$, but $F \succ F\backslash\{f_I\} \cup \{h\}$, which contradicts Axiom 2. ■

Claim 7. *For all* $s \in I \in \sigma(\widehat{p})$, $\pi(s|I) = \frac{\mu(s)}{\mu(I)}$.

Proof. Using Claim 6,

$$\mu(I) := \sum_{s'\in I}\mu(s') = \frac{\mu(s)}{\pi(s|I)}\sum_{s'\in I}\pi(s'|I) = \frac{\mu(s)}{\pi(s|I)}$$
$$\Rightarrow \pi(s|I) = \frac{\mu(s)}{\mu(I)}.$$

■

Define $\rho(I) := \frac{\widehat{p}(I)}{\mu(I)}$. Using Claim 7, we can substitute $\mu(s)\rho(I)$ for $\pi(s|I)\,\widehat{p}(I)$ in (4). Bayes' law implies that $\rho_s(I) = \begin{cases} \rho(I) & \text{if } s \in I \\ 0 & \text{if } s \notin I \end{cases}$.

6.4. Proof of Theorem 2

(if) Let Ψ be a uniform cover of S'. Let $k \geq 1$ be the smallest number of times that S' is covered by Ψ. Set $\rho(I) = \frac{\beta(I)}{k}$ for all $I \in \Psi$.
(only if) Suppose that $\rho : 2^{S'} \rightarrow [0,1]$ is a generalized partition, with $\sigma(\rho) = \Psi$. In addition to $\rho(I) = 0$ for $I \notin \Psi$, the conditions that ρ should satisfy can be written as $\mathbf{A}\rho_\Psi = 1$, where $\mathbf{A}$ is a $|S'| \times |\Psi|$ matrix with entries $a_{i,j} = \left\{ \begin{array}{ll} 1 & s \in I \\ 0 & s \notin I \end{array} \right|$, ρ_Ψ is a $|\Psi|$-dimensional vector with entries $(\rho(I))_{I \in \Psi}$, and $\mathbf{1}$ is a $|S'|$-dimensional vector of ones.

Suppose first that $\rho(I) \in \mathbb{Q} \cap (0,1]$ for all $I \in \Psi$. Rewrite the vector ρ_Ψ by expressing all entries using the smallest common denominator, $\xi \in \mathbb{N}_+$. Then Ψ is a generalized partition of size ξ. To see this, let $\beta(I) := \xi\rho(I)$ for all $I \in \Psi$. Then

$$\sum\nolimits_{I \in \Psi | s \in I} \beta(I) = \sum\nolimits_{I \in \Psi | s \in I} \xi\rho(I) = \xi$$

for all $s \in S'$.

It is thus left to show that if $\rho_\Psi \in (0,1]^{|\Psi|}$ solves $\mathbf{A}\rho_\Psi = \mathbf{1}$, then there is also $\rho'_\Psi \in [\mathbb{Q} \cap (0,1]]^{|\Psi|}$ such that $\mathbf{A}\rho'_\Psi = \mathbf{1}$.

Let $\widehat{P}$ be the set of solutions for the system $\mathbf{A}\rho_\Psi = \mathbf{1}$. Then, there exists $X \in \mathbb{R}^k$ (with $k \leq |\Psi|$) and an affine function $f : X \rightarrow \mathbb{R}^{|\Psi|}$ such that $\widehat{\rho_\Psi} \in \widehat{P}$ implies $\widehat{\rho_\Psi} = f(x)$ for some $x \in X$. We first make the following two observations:

(i) there exists f as above, such that $x \in \mathbb{Q}^k$ implies $f(x) \in \mathbb{Q}^{|\Psi|}$;

(ii) there exists an open set $\widetilde{X} \subseteq \mathbb{R}^k$ such that $f(x) \in \widehat{P}$ for all $x \in \widetilde{X}$

To show (i), apply the Gauss elimination procedure to get f and X as above. Using the assumption that $\mathbf{A}$ has only rational entries, the Gauss elimination procedure (which involves a sequence of elementary operations on $\mathbf{A}$) guarantees that $x \in \mathbb{Q}^k$ implies $f(x) \in \mathbb{Q}^{|\Psi|}$.

To show (ii), suppose first that $\rho^* \in \widehat{P} \cap (0,1)^{|\Psi|}$ and $\rho^*_\Psi \notin \mathbb{Q}^{|\Psi|}$. By construction, $\rho^*_\Psi = f(x^*)$, for some $x^* \in X$. Since $\rho^*_\Psi \in (0,1)^{|\Psi|}$ and f is affine, there exists an open ball $B_\varepsilon(x^*) \subset \mathbb{R}^k$ such that $f(x) \in \widehat{P} \cap (0,1)^{|\Psi|}$ for all $x \in B_\varepsilon(x^*)$, and in particular for $x' \in B_\varepsilon(x^*) \cap \mathbb{Q}^k$ $(\neq \phi)$. Then $\rho'_\Psi = f(x') \in [\mathbb{Q} \cap (0,1]]^{|\Psi|}$. Lastly, suppose that $\rho^*_\Psi \in \widehat{P} \cap (0,1]^{|\Psi|}$ and that there are $0 \leq l \leq |\Psi|$ sets $I \in \Psi$, for which $\rho(I)$ is uniquely determined to be 1. Then set those l values to 1 and repeat the above procedure for the remaining system of $|\Psi| - l$ linear equations.

6.5. Proof of Theorem 3

It is easy to check that any preferences with a subjective filtration representation as in Theorem 3 satisfy Axiom 3. The rest of the axioms are satisfied since Theorem 3 is a special case of Theorem 1.

To show sufficiency, first observe that by Axiom 3 and Claim 2, I,$I' \in \sigma(\rho)$ implies that either $I \subset I'$, or $I' \subset I$, or $I \cap I' = \emptyset$. This guarantees that for any $M \subset \sigma(\rho)$ and $s \in \sigma(\mu)$, $\arg\max_{I \in M} \{|I| \,|s \in I\}$ is unique if it exists.

For any state $s \in \sigma(\mu)$, let $I_1^s = \arg\max_{I \in \sigma(\rho)} \{|I| \,|s \in I\}$. Define $T_1 := \{I_1^s \,|s \in \sigma(\mu)\}$. Let $\eta_1 = \min_{I \in T^1} (\rho(I))$. Set

$$\rho_1(I) = \begin{cases} \rho(I) - \eta_1 & \text{if } I \in T_1 \\ \rho(I) & \text{if } I \notin T_1 \end{cases}.$$

Let $\rho_n : \sigma(\rho) \to [0,1]$ for $n \in \mathbb{N}$. Inductively, if for all $s \in \sigma(\mu)$ there exists $I \in \sigma(\rho_n)$ such that $s \in I$, then for any $s \in \sigma(\mu)$ let $I_{n+1}^s = \arg\max_{I \in \sigma(\rho_n)} \{|I| \,|s \in I\}$. Define $T_{n+1} := \left\{I_{n+1}^s \,|s \in \sigma(\mu)\right\}$. Let $\eta_{n+1} = \min_{I \in T_{n+1}} (\rho_n(I))$. Set

$$\rho_{n+1}(I) = \begin{cases} \rho_n(I) - \eta_{n+1} & \text{if } I \in T_{n+1} \\ \rho_n(I) & \text{if } I \notin T_{n+1} \end{cases}.$$

Let $N+1$ be the first iteration in which there exists $s \in \sigma(\mu)$ which is not included in any $I \in \sigma(\rho_N)$. Axiom 1 implies that N is finite and that $(T^n)_{n=1,..,N}$ is a sequence of increasingly finer partitions, that is, for $m > n$, $I_m^s \subseteq I_n^s$ for all s, with strict inclusion for some s.

Claim 8. *$\rho(I) = \sum_{n \leq N | I \in T_n} \eta_n$ for all $I \in \sigma(\rho)$.*

Proof. First note that by the definition of N, $\rho(I) \geq \sum_{n \leq N | I \in T_n} \eta_n$ for all $I \in \sigma(\rho)$. If the claim were not true, then there would exist $I' \in \sigma(\rho)$ such that $\rho(I') > \sum_{n \leq N | I' \in T_n} \eta_n$. Pick $s' \in I'$. At the same time, by the definition of N, there exists $s'' \in \sigma(\mu)$ such that if $s'' \in I \in \sigma(\rho)$ then $\rho(I) = \sum_{n \leq N | I \in T_n} \eta_n$. We have,

$$\begin{aligned} \mu(s'') &= \textstyle\sum_{I \in \sigma(\rho)} \Pr(s'' \,|I) \,\rho(I) \,\mu(I) = \sum_{I \in \sigma(\rho)} \Pr(s'' \,|I) \,\mu(I) \sum_{n \leq N | I \in T_n} \eta_n \\ &= \textstyle\sum_{n \leq N} \Pr\left(s'' \middle| I_n^{s''}\right) \mu\left(I_n^{s''}\right) \eta_n = \mu(s'') \sum_{n \leq N} \eta_n, \end{aligned}$$

where the last equality follows from Claim 7. Therefore, $\sum_{n \leq N} \eta_n = 1$. At the same time

$$\begin{aligned} \mu(s') &= \textstyle\sum_{I \in \sigma(\rho)} \Pr(s' \,|I) \,\rho(I) \,\mu(I) > \sum_{I \in \sigma(\rho)} \Pr(s' \,|I) \,\mu(I) \sum_{n \leq N | I \in T_n} \eta_n \\ &= \textstyle\sum_{n \leq N} \Pr\left(s' \middle| I_n^{s'}\right) \mu\left(I_n^{s'}\right) \eta_n = \mu(s') \sum_{n \leq N} \eta_n = \mu(s'), \end{aligned}$$

which is a contradiction. ■

Claim 8 implies that $\sigma(\rho_{N+1}) = \emptyset$. Let $\eta_m := 0$ and for $t \in [0,1)$ define the filtration $\{\mathcal{P}_t\}$ by

$$\mathcal{P}_t := T_n, \text{ for } n \text{ such that } \textstyle\sum_{m=0}^{n-1} \eta_m \le t < \sum_{m=0}^{n} \eta_m.$$

The pair $(\mu, \{\mathcal{P}_t\})$ is thus a subjective filtration. The next claim establishes that $(\mu, \{\mathcal{P}_t\})$ is unique.

Claim 9. *If $\left(\widehat{\mu}, \left\{\widehat{\mathcal{P}}_t\right\}\right)$ induces a representation as in Theorem 3, then $\left(\widehat{\mu}, \left\{\widehat{\mathcal{P}}_t\right\}\right) = (\mu, \{\mathcal{P}_t\})$.*

Proof. μ and ρ are unique according to Theorem 1. We already observed that $I \cap \widehat{I} = \emptyset$, $\widehat{I} \subset I$, or $I \subset \widehat{I}$ for any $I, \widehat{I} \in \sigma(\rho)$. Suppose that $\{\mathcal{P}_t\} \neq \left\{\widehat{\mathcal{P}}_t\right\}$. Then there exist $t \in (0,1)$ and $I \in \sigma(\rho)$, such that $I \in \mathcal{P}_t$ and $I \notin \widehat{\mathcal{P}}_t$. Fix $s \in I$. There is $\widehat{I} \in \widehat{\mathcal{P}}_t$ with $s \in \widehat{I}$ and, therefore, either $\widehat{I} \subset I$ or $I \subset \widehat{I}$. Assume, without loss of generality, that $\widehat{I} \subset I$. Let $M = \{I' \in \sigma(\rho) : I \subseteq I'\}$. Let $\rho(M) := \sum_{I \subset M} \rho(I)$ and $\mu(M) = \sum_{I \subset M} \mu(I) = \sum_{I \subset M} \sum_{s \in I} \mu(s)$. Then according to $(\mu, \{\mathcal{P}_t\})$, $\rho(M)\mu(M) \ge t$, while according to $\left(\widehat{\mu}, \left\{\widehat{\mathcal{P}}_t\right\}\right)$, $\rho(M)\mu(M) < t$, which is a contradiction to the uniqueness of (μ, ρ) in Theorem 1. ■

The last claim formalizes the observation in Remark 2.

Claim 10. *If $\succeq$ also satisfies Axiom 4, then $\mathcal{P}_0 = \{\sigma(\mu)\}$.*

Proof. Suppose to the contrary, that there are $\{I, I'\} \subset \mathcal{P}_0$ such that $I \cap I' = \emptyset$ and $I \cup I' \subseteq \sigma(\mu)$. Then, any saturated F includes some act h with $\sigma(h) \subset I$ and another act g with $\sigma(g) \subset I'$, but it does not include an act that contains both h and g, which contradicts Axiom 4. ■

6.6. Proof of Theorem 4

Let $\succeq$ be represented as in Theorem 1. Consider the menu $\{c, f\}$. We make the following two observations: first, $\{f'\} \sim \{c\}$ if and only if $s \notin \sigma(\mu)$. Second, since conditional on any $I \ni s, s'$

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

and since $\{c\} \succ \{f\}$, $\sum_{\widehat{s} \in I} f(\widehat{s})\,\mu(\widehat{s}) > c \sum_{\widehat{s} \in I} \mu(\widehat{s})$ if and only if $s \in I$ but $s' \notin I$. These are the only events in which DM expects to choose f from $\{c, f\}$. Note that

$$V(\{c, f\}) = c + \mu(s)\,k \textstyle\sum_{I | s \in I,\ s' \notin I} \rho(I)$$

and that

$$V(\{f'\}) = c + \mu(s)\,k'.$$

Therefore, $\{f'\} \succeq \{c, f\}$ if and only if $\sum_{I|s\in I,\ s'\notin I} \rho(I) \leq \frac{k'}{k}$.

By the first observation, Definition 12 (i) is equivalent to the condition $\sigma(\mu^1) = \sigma(\mu^2)$. By the second observation, Definition 12 (ii) is equivalent to the condition

$$\sum_{I|s\in I,\ s'\notin I} \rho^1(I) \leq \frac{k'}{k} \Rightarrow \sum_{I|s\in I,\ s'\notin I} \rho^2(I) \leq \frac{k'}{k}$$

for all $k' \in [0, k]$, or,

$$\sum_{I|s\in I,\ s'\notin I} \rho^1(I) \geq \sum_{I|s\in I,\ s'\notin I} \rho^2(I).$$

6.7. Proof of Theorem 5

(i) DM1 does not learn earlier than DM2 $\Leftrightarrow$

there exists t such that $\mathcal{P}_t^1$ is not finer than $\mathcal{P}_t^2$ $\Leftrightarrow$

there exists two states s, s', such that $s, s' \in I$ for some $I \in \mathcal{P}_t^1$, but $s, s' \notin I'$ for any $I' \in \mathcal{P}_t^2$ $\Leftrightarrow$

$\Pr^2(\{I\,|s \in I,\ s' \notin I\}\,|s) = \sum_{I|s\in I, s'\notin I} \rho^2(I) \geq 1 - t$, but $\Pr^1(\{I\,|s \in I,\ s' \notin I\}\,|s) = \sum_{I|s\in I, s'\notin I} \rho^1(I) < 1 - t$ $\Leftrightarrow$

DM1 does not value binary bets more than DM2.

(ii) **(if)** Suppose $\{\mathcal{P}_t^1\}$ is weakly finer than $\{\mathcal{P}_t^2\}$ and that $\mu^1 = \mu^2 = \mu$. Fix a time t. Any $I \in \mathcal{P}_t^2$ is measurable in $\mathcal{P}_t^1$, that is, there is a collection of sets $I_k \subset \mathcal{P}_t^1$ such that $I = \underset{k}{\cup} I_k$. Since the max operator is convex, $\sum_{I_k} \max_{f\in F}\left[\sum_{s\in I_k} f(s)\mu(s)\right] \geq \max_{f\in F}\left[\sum_{s\in I} f(s)\mu(s)\right]$. Since t was arbitrary, we have

$$\begin{aligned} V^1(F) &= \int_{[0,1]} \left\{\sum_{I\in\mathcal{P}_t^1} \max_{f\in F}\left[\sum_{s\in I} f(s)\mu(s)\right]\right\} dt \\ &\geq \int_{[0,1]} \left\{\sum_{I\in\mathcal{P}_t^2} \max_{f\in F}\left[\sum_{s\in I} f(s)\mu(s)\right]\right\} dt = V^2(F). \end{aligned}$$

Claim 1 in DLST states that if DM1 has mor preference for flexibility than DM2 then $\{f\} \succeq_1 \{g\}$ if and only if $\{f\} \succeq_2 \{g\}$, which means that for any f, $V^1(\{f\}) = V^2(\{f\})$. Therefore, $V^2(F) \geq V^2(\{f\})$ implies $V^1(F) \geq V^1(\{f\})$.

(only if) By Theorem 2 in DLTS, more preference for flexibility implies that $\mu^1 = \mu^2$. It is left to show that having more preference for flexibility implies learning earlier. For $i = 1, 2$, let

$$t^i(I) = \min\left\{t\,\middle|\, I \text{ is measurable in } \mathcal{P}_t^i\right\}$$

if defined, otherwise let $t^i(I) = 1$. Let,

$$\Delta^i(I) = \max\left\{t \,\middle|\, I \in \mathcal{P}_t^i\right\} - \min\left\{t \,\middle|\, I \in \mathcal{P}_t^i\right\}$$

if defined, otherwise let $\Delta^i(I) = 0$. We make the following intermediate claim.

Claim 11. *DM1 has more preference for flexibility than DM2 implies that for all $I \in 2^{\sigma(\mu)}$*

$$\textstyle\sum_{I' \subseteq I} \Delta^1(I')\,\mu(I') \geq \sum_{I' \subseteq I} \Delta^2(I')\,\mu(I').$$

Proof. Suppose that there is $I \in 2^S$ with $\sum_{I' \subseteq I} \Delta^2(I')\,\mu(I') > \sum_{I' \subseteq I} \Delta^1(I')\,\mu(I')$. Obviously $I \subsetneq \sigma(\mu)$. Define the act f by

$$f(s) = \begin{cases} \delta > 0 \text{ if } s \in I \\ 0 \text{ if } s \notin I \end{cases}.$$

Let c denote the constant act that gives $c > 0$ in every state, such that $\delta > c > \frac{\mu(I)}{\mu(I'')}\delta$ for all $I'' \in 2^{\sigma(\mu)}$ with $I \subsetneq I''$. Then $V_i(\{f, c\}) = c + (\delta - c)\sum_{I' \subseteq I} \Delta^i(I')\,\mu(I')$. Finally, pick c' such that

$$\textstyle(\delta - c)\sum_{I' \subseteq I} \Delta^2(I')\,\mu(I') > c' - c > (\delta - c)\sum_{I' \subseteq I} \Delta^1(I')\,\mu(I'),$$

to find $\{f, c\} \succ_2 \{c'\}$ but $\{c'\} \succ_1 \{f, c\}$, and hence DM1 cannot have more preference for flexibility than DM2. ■

Under the assumptions of Theorem 3,

$$\textstyle\sum_{I' \subseteq I} \mu^i(I')\,\Delta^i(I') = \mu^i(I)\left(1 - t^i(I)\right).$$

By Claim 11, DM1 has more preference for flexibility than DM2 implies that $t^1(I) \leq t^2(I)$ for all I, which is equivalent to $\{\mathcal{P}_t^1\}$ being weakly finer than $\{\mathcal{P}_t^2\}$.

References

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

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

[3] Dekel, E., B. Lipman, and A. Rustichini (2009), "Temptation-driven preferences," *Review of Economic Studies*, 76 (3), 937–971.

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

[5] Dillenberger, D., J.S. Lleras, P. Sadowski, and N. Takeoka (2012), "A Theory of Subjective Learning," PIER Working Paper number 12-034.

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

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

[8] Gilboa, I. (2009), " Theory of Decision under Uncertainty," Cambridge University Press.

[9] Gilboa, I., A. Postlewaite, and D. Schmeidler (2009), "Is It Always Rational to Satisfy Savage's Axioms?," *Economics and Philosophy*, 25, 285–296.

[10] Gilboa, I., A. Postlewaite, and D. Schmeidler (2009), "Rationality of Belief Or: Why Savage's axioms are neither necessary nor sufficient for rationality," *Synthese*, forthcoming.

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

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

[13] Lehrer, E. (2012), "Partially Specified Probabilities: Decisions and Games," *American Economic Journal: Microeconomics*, Vol. 4, No. 1, 70-100.

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

[15] Riella, G. (2011), "Preference for Flexibility and Dynamic Consistency," mimeo.

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

[17] Shapley, L. S. (1967), "On Balanced Sets and Cores." *Naval Research Logistics Quarterly,* 14, 453–460.

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