

Incomplete Menu Preferences and Ambiguity ^{*}

Madhav Chandrasekher[†]

Current Draft: January 21, 2012

Abstract: This paper models decision problems where the decision-maker (DM) may not be able to compare two risky acts. Incomparability of acts can, in principle, stem from two sources. First, perception of risk is subjective and the DM might have only a coarse assessment of the probability measure over the space of ex post outcomes that determine the values of the acts. Second, there may be co-dependence on the realized outcome and the set of the consumption choices that are available conditional on that outcome. Existing models of incomplete preferences, which build on Bewley's (1986) model of Knightian Uncertainty, accommodate the first consideration, but foreclose the second possibility as a source of incompleteness. This paper presents an axiomatic model of incomplete preferences that accounts for both of these possibilities as a source of incompleteness.

Keywords: Menu Choice, Incomplete Preferences, Knightian Uncertainty.

^{*}This paper has benefitted from the comments of the audiences of the Spring 2010 Midwest Theory meeting at Northwestern, SWET 2011 at CalPoly, San Luis Obispo, and RUD 2012, at the Collegio Carlo Alberto. I am also grateful to Simone Cerreia-Vioglio, Chris Chambers, Juan Dubra, Edi Karni, Peter Klibanoff, Fabio Maccheroni, Efe Ok, Luca Rigotti, Philipp Sadowski, and Kyoungwon Seo for helpful comments. This paper was completed while I was a visitor at the Math Center of the Kellogg School at Northwestern University. I thank them for their hospitality. Comments are welcome.

[†]Mailing address: Department of Economics, W.P. Carey School of Business, P.O. Box 873806, Tempe, AZ 85287-3806; E-mail address: madhav.chandrasekher@asu.edu

1 Introduction

This paper generalizes Bewley’s classic formulation of Knightian ambiguity to a setting where the decision-maker’s domain of choice is the set of subsets of $\Delta(X)$, the space of lotteries on a finite set X . Preferences over this domain have been given the moniker “preferences over subjective contingencies” after the seminal paper by DLR (2001) that analyzed such preferences and provided an important representation result. The DLR paper shows that preferences on $2^{\Delta(X)}$ that satisfy the Herstein-Milnor mixture space axioms admit a utility representation that takes the form of a (state-dependent) expected utility functional. Moreover, the utility kernel and the probability space over which the expectation operator is defined are shown to be essentially unique.¹

The DLR Theorem provides a beautiful counterpart (on both the axiomatic and functional form side) of the Anscombe-Aumann Theorem. The model interprets a menu as an implicit Anscombe-Aumann act, where the domain of the act (the subjective state space) is in the mind of the decision-maker. An important axiom implied by the DLR representation is the order axiom, which requires preferences over menus to be complete. The starting point of this paper is to analyze what happens when we relax this assumption. Incomplete preferences were famously analyzed by Bewley (1986) who provided the prevailing utility representation for incomplete preferences. Bewley’s representation considers an Anscombe-Aumann setup, maintaining the same axiomatic framework except for the order axiom on compound Anscombe-Aumann acts- so that preferences are allowed to be incomplete. Bewley’s representation theorem suggests that the cause of incompleteness is that the decision-maker (DM) has imprecise probabilistic beliefs over the (objective) state-space over which acts are defined. Thus, rather than a single probability measure over the state space the representation posits an entire set of measures - so that the issue of which among two acts is preferred depends on the particular belief used to compute the subjective expected utility of that act.

This is a powerful result that has, to date, provided the prevailing paradigm to ‘visualize’ preference incompleteness.² The primary conceptual goal of the current project to further extend our understanding of why preferences are incomplete. Our model also uses the Bewley criterion for comparing menus; however, the choice domain is different from the Bewley framework - the behavioral primitive in this paper is an incomplete preference order over menus. Our main result provides an

¹This requires fixing a normalization of the state-space and state-dependent utility kernel at the outset.

²To be clear, there have been many axiomatic models of incomplete preferences in the decision theory literature - a handful of which are listed in the references section. However, a “multiple-utility” criterion is a common feature of all these models.

axiomatic foundation for a Bewley-type criterion for representing incomplete preference relations on menus.

As a motivating example, consider a college bound senior who is choosing among two choices, say college A and college B. This might be a non-trivial choice problem if the DM has to commit to a college choice before knowing with certainty what her consumption preferences (over the amenities offered by various colleges) will be. Let us use the Anscombe-Aumann machinery to model this decision problem. For simplicity assume the space of uncertainty facing the DM is objective so that any two decision-makers faced with the same college choice problems have the same space of uncertainty - for example, the space of all possible rankings of career options. A college can then be coded as an Anscombe-Aumann act as illustrated below. Take the state space to be the set of rankings over career options available to the DM, and for a parsimonious description of the act associated with a college identify a state with its top-ranked career choice.

$$\left[\begin{array}{cc} \text{doctor} & \text{biology} \\ \text{physicist} & \text{physics} \\ \vdots & \vdots \\ \text{veterinarian} & \text{veterinary science} \end{array} \right]$$

In the above example, a college is coded into an act as follows. For each state, the DM picks the major at that college that puts him on track towards the most preferred career choice (in that state). Once we code colleges into acts using this procedure, the question turns to how the DM assigns utilities to these acts. The model in this paper lies within the subjective expected utility (SEU) paradigm, so that for each belief that the DM has about the distribution of states (she may have several) she computes the value of an act by integrating against each subjective belief over the distribution over states. This yields a collection of utilities for each menu-induced act. Using the Bewley criterion, we say that college A is preferred to college B if the prior by prior SEU value of the act induced by college A dominates that of the act induced by college B.

So far, this description sounds just like the standard, i.e. objective Anscombe-Aumann, Bewley model. While the utility models in the objective and subjective (i.e. menus) setting are similar, there is an important issue that arises in the menu choice environment that is absent in the objective setting. A key assumption in applying the objective Anscombe-Aumann framework is that the set of feasible ex post choices is *independent* of the realized ex post state. This assumption is built into the behavioral primitive, which is a relation on lottery-valued acts whose domain is the objective state space. Since the range of these acts is the full set of vNM lotteries, this implicitly makes two assumptions about the nature of the choice

problem. First, we can arbitrarily perturb the ex post payoff for any given act and remain within the set of feasible acts. Second, the set of all these perturbed acts is in the support of the preference relation. Note that in the college choice example both of these restrictions can be violated. For example, liberal arts colleges might offer a different set of academic options than engineering schools so that when we encode colleges as (implicit) Anscombe-Aumann acts, the set of menu-induced acts is much smaller than the set of objective Anscombe-Aumann acts. Thus, when there is co-dependence between the realized state and the consumption choices available in that state, the domain of the preference relation is much smaller.

It is particularly important to take the co-dependence between ex post states and ex post choices into account when trying to understand why preferences may be incomplete. Indeed, one of the very sources of incompleteness may be the presence of this co-dependence.³ For example, the reason I cannot compare A and B might simply be because they offer different consumption options in different states - an explanation that is not obviously rooted in my ambiguous beliefs about the distribution of ex post states. This can be an additional source of incompleteness since there are fewer chains of comparisons that the DM can make in comparing two objects. For example, in the objective set-up a DM can compare two acts (f, g) by successively perturbing the payoffs of the initial act f and finding a sequence of perturbations $f_0 = f, f_1, \dots, f_n = g$ such that $f = f_0 \succeq f_1 \succeq \dots \succeq f_n = g$. Transitivity then implies that $f \succeq g$. When the support of the preference relation is not rich enough, these intermediate comparisons may not be available, thus (possibly) leading to an additional source of incompleteness.

Thus, in the (subjective) menu choice framework there are possibly multiple sources of incompleteness. On the one hand, there might be objective ambiguity about the governing distribution of ex post states - as in the objective set-up. On the other hand, there is co-dependence between the realized state and the consumption choices available in that state - so that the support of the preference relation is not rich enough to resolve comparisons of pairs (f, g) via a sequence of comparisons among perturbed acts. The main result in the paper shows that the menu-choice analogue of the Bewley axioms are necessary and sufficient for a representation of the decision procedure we informally discussed above. On a conceptual level, the value of the result is that it extends the analysis of incomplete preferences to the domain of menu choice. Moreover, the representation ‘visualizes’ preference incompleteness through the same mechanism as in the classical Bewley model, i.e. through ambiguity over the distribution of ex post states. From an axiomatic viewpoint, we

³The question of “sources of incompleteness” is the topic of the paper Ok-Ortoleva-Riella (2009). Thus, we are not the first to seek out sources of incompleteness other than Knightian uncertainty. However, to our knowledge, we are the first to investigate whether the sparseness of the choice domain is a source of incompleteness.

find this a surprising result. Our axioms are essentially a menu-choice translation of the axioms in the Bewley framework, yet the domain of the preference relation is much smaller. As discussed above, this suggests that *ceteris paribus* the sparseness of the choice domain should compound whatever pre-existing incompleteness in preferences is already present due to ambiguity. The representation result shows that while our behavioral primitive accommodates an additional source of incompleteness, the DM, nevertheless, behaves as if the sole source of incompleteness is ambiguity about the distribution on the state space. Our proof of this result is novel and uses techniques that are distinct from those typically used in the ambiguity literature. As I comment in further detail below, standard techniques encounter some initial difficulties, which is the reason we use an alternative approach.

A standard technique applied in the objective ambiguity framework is to first make an implicit identification between AA acts and utility vectors in $\mathbf{R}^{|S|}$ (where S is the objective state space). The set of measures that capture the DM's ambiguity is then recovered from the dual cone to the cone in utility space generated by the set of differences $f - g$, where $f \succeq g$ (call this the *pointed cone*). A major impediment in carrying out this procedure in our setting is the sparseness of the support in the preference relation. Specifically, a technical hypothesis that is used to apply the duality between the pointed cone and dual cone is that the algebraic interior of the set of induced utility vectors (in its span) is non-empty. This condition is trivially satisfied in the objective AA set-up precisely because the set of ex post choices is the *full* domain of vNM lotteries, and hence independent of the realized state. Thus, for *every* utility vector that is the image of an AA-act we can find an appropriately scaled basis for a positive orthant with this vector at the origin such that every point in the orthant is a multiple of the image of an AA-act.

This latter property is quite strong and much more than what is needed in order to apply the duality between the pointed cone and its dual. Moreover, it is seems unlikely (though we have not proven otherwise) that this property holds for the menu choice domain. Thus, while our method also exploits the link between a pointed cone and its dual, the verification step in applying this duality (checking the non-emptiness of the algebraic interior) is a non-trivial part in the argument.⁴

⁴One may be able to obtain non-emptiness of the algebraic interior by standard approaches - so that we know that the pointed cone is closed. However, this will only yield a set of possibly signed measures whose associated hyperplanes support the cone. In the Anscombe-Aumann framework one applies the monotonicity axiom to show that these measures must in fact be probability measures. This works precisely because the monotonicity axiom implies that the pointed cone has non-empty interior (in the ambient vector space). In the menus setting, showing that the pointed cone has non-empty interior is non-trivial and constitutes the heart of our argument. I thank Fabio Maccheroni for a helpful discussion on this point.

The representation result yields, as a corollary, a subjective Anscombe-Aumann representation theorem, where the domain of choice is menus rather than objective AA acts. This result is disjoint from the DLR (2001) representation (which we also interpret as an Anscombe-Aumann-type representation) in the sense that neither model nests the other.

I have commented above on why the standard Shapley-Baucells argument does not immediately apply to our setting. However, it is also natural to ask whether we can apply the DLR (2001) technology in our setting. Let us give a quick synopsis of the DLR strategy.⁵ First, obtain a linear functional on the space support functions that represents the given preference on menus.⁶ Now use Hahn-Banach to lift the linear functional on this subspace to a linear functional on the full space of functions within which the support functions are embedded. Finally, apply continuity of the functional (and duality) to represent the functional as integration against a (possibly signed) measure on the state space (which is the domain of the support function).

There are two obstructions to carrying out this recipe in our setting. First, the integrand in our utility representation is not a support function, but is instead a *Strotzian value function* associated to a menu. In words, for a given menu the associated Strotzian value function maps a state to the indirect utility obtained from the given menu in that state, where the indirect utility is computed using the well-known Strotz utility formula (e.g. see Gul-Pesendorfer (GP) (2005)). The problem is then the following: While the map $M \mapsto \sigma_M$ is a homeomorphism (here σ_M is a support function), the map $M \mapsto \phi_M$ is *not* a homeomorphism (here ϕ_M is the Strotzian value function associated to the menu). Thus, linear menu preferences that satisfy Hausdorff continuity do not map homeomorphically into the space of Strotzian value functions.

One can try to remedy this situation in the following manner. First, find a pair of topologies (τ_1, τ_2) on (resp.) the space of menus and the space of Strotzian value functions such that menus topologized under τ_1 map homeomorphically to Strotzian functions topologized under τ_2 .⁷ Second, prove that the space of Strotzian value functions is dense (in the topology τ_2) subspace of $C(\mathcal{S})$ - the space of all continuous functions . Third, prove a version of the Riesz representation theorem

⁵The remainder of this introduction can be skipped by readers willing to take on faith that our main result cannot be obtained by a modification of the DLR arguments. The only purpose of this discussion is to convince a reader well-versed in the DLR argument that non-trivial complications arise when one attempts to import the DLR argument to the model of this paper.

⁶Recall that (i) a menu is identified with the support function induced by the menu, (ii) this correspondence between (closed and convex)) menus and support functions is bijective.

⁷We actually only need τ_2 since we can just obtain τ_1 by “pulling-back” open sets via the map that takes a menu to its Strotzian value function.

for the space $(C(\mathcal{S}), \tau_2)$ to decompose the lifted representation as integration against a measure on \mathcal{S} . The main difficulty is in the last two steps.

I can not speak to the feasibility of the third step (generalizing the Riesz theorem) since that is not the path I follow in this paper. However, it is more transparent to see that the density argument in DLR does not apply. Recall that a key step of their density argument invoked the lattice structure of the set of differences, $\{\sigma_1 - \sigma_2 : \sigma_i \geq 0\}$, where the σ_i are elements of the cone of non-negative support functions. This lattice structure does not exist when we replace support functions with Strotzian value functions.⁸ This summarizes the main difficulty in trying to import the DLR proof to our setting. As a consequence, we develop a different argument to obtain our representation result. A virtue of the approach in this paper is that it is self-contained.

The structure of this paper is as follows. In the upcoming section we provide a description of the DLR model and, for our purposes, the main results in DLR (2001). We then describe the required modification of these axioms which yield our main result, a generalization of Bewley's theorem. Section 3 is devoted to a proof of this result and section 4 concludes.

2 The DLR Model

Let $X := \{x_1, \dots, x_k\}$ be a finite set and put $\Delta(X) := \{(p_1, \dots, p_k) : \sum_{i=1}^k p_i = 1, 0 \leq p_i \leq 1\}$ the set of lotteries on the set X , viewed as a metric space with the topology induced by the usual Euclidean metric. Put $M(X) := \{M : M \subseteq \Delta(X), M \text{ closed}\}$ and let $d(\cdot, \cdot)$ be the metric induced by the Hausdorff topology (induced by the Euclidean metric). Consider a preference relation (referred to as DLR preferences) on $M(X)$ defined by the following axioms:

A0: (Non-Triviality) $\exists M, M' \in M(X)$ such that $M \succ M'$.⁹

A1': (Order) $(\succeq, M(X))$ is a complete, rational order.

A2: For each menu B , the sets $\{A \in M(X) : A \succeq B\}$ and $\{A \in M(X) : B \succeq A\}$

⁸To see why this is not the case note that the supremum of two Strotzian value functions is not necessarily a Strotzian value function.

⁹The only role of A0 is to ensure uniqueness in the utility representation. Indeed, for our model the indeterminacy is quite stark. Without A0, the representation is completely indeterminate - any set of measures can be used to represent the trivial menu preference where all menus are indifferent. On the other hand, with A0 the representing set of measures is unique.

are closed.¹⁰

A3: (Independence) If $A \succeq B$, then $\forall \alpha \in (0, 1)$ and any $C \in M(X)$, $\alpha \cdot A + (1 - \alpha) \cdot C \succeq \alpha \cdot B + (1 - \alpha) \cdot C$, where we put $\alpha \cdot A + (1 - \alpha) \cdot C := \{l \in \Delta(X) : l = \alpha p + (1 - \alpha)q, p, q \in A, C \text{ resp.}\}$

Put $S := \{(u_1, \dots, u_k) \in \mathbf{R}_+^k : \sum_{i=1}^k u_i = 1\}$ and for each $s \in S$ define a vNM utility on $\Delta(X)$ by the formula $U_s(p_1, \dots, p_k) = \sum_{i=1}^k p_i u_i$. Clearly any given vNM preference is identified by an equivalence class of affine utilities. It is a simple exercise to check that for any affine utility $U(\cdot)$, there are constants a, b with $a > 0$ such that the utility $U' = aU + b$ has indices $u_i \geq 0$ and $\sum_{i=1}^k u_i = 1$. This process does *not* lead to a unique element of \mathcal{S} , so that \mathcal{S} is a superset of the set of vNM preferences over $\Delta(X)$. Contrast this with the space that DLR (2001) use as a representative set of vNM preferences,

$$\mathcal{S}_{\text{DLR}} := \{s = (u_1, \dots, u_k) \in \mathbf{R}^k : (i) \sum_{i=1}^k u_i = 0, (ii) \sum_{i=1}^k u_i^2 = 1\}$$

These two conditions on the vectors (u_1, \dots, u_k) assure that if $u, u' \in \mathcal{S}$ and $u = a \cdot u' + b$, then $a = 1, b = 0$. Letting Π_{vNM} denote the set of vNM preferences we can explicitly identify Π_{vNM} with a quotient of \mathcal{S} and endow Π_{vNM} with the corresponding quotient topology. To see why these remarks are relevant recall the basic DLR strategy. Let $\Sigma := \{M \subseteq \Delta(X) : M \text{ convex, closed}\}$ and following DLR, define a map $\Theta : \Sigma \hookrightarrow C(\mathcal{S}_{\text{DLR}})$ by

$$\Theta(M)(s) := \max_{x \in M} U_s(x)$$

An important consequence of this is the following:

Fact: Θ induces an affine homeomorphism between $M(X)$ and the subset of $C(\mathcal{S}_{\text{DLR}})$ defined by $\Theta(M(X))$.

The fact allows DLR to transfer the analysis of preferences on the space $M(X)$ to preferences on $C(\mathcal{S}_{\text{DLR}})$, by identifying a menu with the vector of ex-post utilities induced by this menu. The key result in DLR (2001) for our purposes is:

Theorem 1. (DLR) Let $(\succeq, M(X))$ be a preference relation satisfying A0 – A3. Then, there is a subset $S_0 \subseteq S$ and a signed measure π (on S) with $\text{supp}(\pi) = S_0$ such that \succeq has the following characterization: $M \succeq M' \Leftrightarrow U(M) \geq U(M')$, where $U(M) := \int_S \Theta(M)(s) \pi_s$.

¹⁰In the corrigendum to DLR (2001), it was noted that this form of continuity was not sufficient to imply the DLR representation theorem. For this paper, however, we will only require a weaker notion of continuity that is more closely related to Hausdorff continuity than the L -continuity axiom required by the general DLR model, hence our sloppiness in writing down the DLR axioms.

The result says that we can think of the choice of a menu M over M' as if it is based on comparing the expected values these menus yield when viewed as elements of $C(\mathcal{S}_{\text{DLR}})$. Since \mathcal{S}_{DLR} can be identified with Π_{DLR} the result offers the nice interpretation that the agent computes the value of a menu by taking a weighted combination of the ex-post values the menu yields in each conceivable ex-post state (i.e. realized vNM preference). If we replace \mathcal{S}_{DLR} with \mathcal{S} in the statement of the theorem we lose this interpretation since Π_{vNM} is a proper quotient of the space \mathcal{S} . However, this problem can be resolved if there is a natural way to project the representation on the space \mathcal{S} onto the space Π_{vNM} such that (i) the measures (or weights) that with domain \mathcal{S} can be pushed forward to the quotient Π_{vNM} , (ii) the utility kernel $U(M) \in C(\mathcal{S})$ can be also be pushed forward to a function $U(M) \in C(\Pi_{\text{vNM}})$, and (iii) the integral on the pushed forward space also generates a cardinal representation of the preference.

Despite the fact that the set Π_{vNM} is a proper quotient of \mathcal{S} , we stick with \mathcal{S} as the operative ‘state space’ in our analysis in this paper. The main reason we do this is twofold. First, since the space \mathcal{S} is convex¹¹ this allows us to push through our proof strategy for the desired representation results (note that the space \mathcal{S}_{DLR} is not convex). Second since the set of vNM preferences is a natural quotient of \mathcal{S} , we use the “push-forward” procedure described in the preceding paragraph to transfer representation results on the space \mathcal{S} onto the real state space of interest, Π_{vNM} . For a given menu M and state s we let $M(s) := \arg \max_{x \in M_s} u(x)$, where $u(\cdot)$ is a fixed vNM representation of the ranking on singleton menus and $M_s := \arg \max_{x \in M} u_s(x)$ (here $u_s(\cdot)$ is a representation of the vNM preference that defines state s).

To the original DLR axioms $A0 - A3$ we add the following:

A4: (monotonicity) Let $M, M' \in M(X)$. If $M(s) \succeq M'(s), \forall s \in S$, then $M \succeq M'$.

To incorporate ambiguity into the model we also require a weakening of $A1'$:

A1: (partial order) $(\succeq, \Delta(X))$ is a total order and $(\succeq, M(X))$ is a (rational) partial order

Finally, the standard continuity axiom ($A2$ as stated above) does not hold for the subjective Bewley model.¹² To formulate the continuity axiom implied by the Bewley model, introduce the following notation. For each menu M define a function

¹¹Note that \mathcal{S}_{DLR} is not convex.

¹²The easiest way to see this is to note that when we impose completeness of the menu order, the Bewley model can be interpreted as the *random Strotz* model of Dekel-Lipman (2007) - which is known to violate the continuity axiom.

$\phi_M : \mathcal{S} \rightarrow \mathbf{R}$ via the formula

$$\phi_M(s) := \max_{f \in M_s} u(f)$$

where $M_s := \arg \max_{f \in M} u_s(f)$. Notice that any two menus M, M' for which $\phi_M = \phi_{M'}$ are equivalent under the menu preference \succeq , by axiom A4. Identify functions ϕ_M with elements of the product space $\mathbf{R}^{\mathcal{S}}$, endowed with the product topology.

To define the continuity axiom introduce the following terminology. A partition, $\{E_i\}$, of \mathcal{S} is called a *Borel partition* if it is generated by a finite collection of hyperplanes $\{s : p_i s \geq c_i\}$. Furthermore, call the Borel partition *coherent* if the constants $c_i \equiv 1/k$. We will shortly see the value of the coherence restriction. For the moment, we restrict attention to the collection of all coherent Borel partitions of \mathcal{S} .¹³ Denote a Borel partition as $\{E_i\}$ and let Σ_{E_i} denote the set of menus M such that ϕ_M is a $\{E_i\}$ -measurable function (call a menu $\{E_i\}$ -measurable when this happens). Let Λ be an index of the set of all coherent Borel partitions (denote a generic partition as $\{E_\alpha\}$) and put

$$\mathcal{A} := \bigcup_{\alpha \in \Lambda} \{E_\alpha\}$$

Abusing notation, let $\Sigma_{\mathcal{A}}$ be the set of all menus whose associated value functions, ϕ_M , are \mathcal{A} -measurable (i.e. $\sigma(\phi_M^{-1}(B) | \forall B \in \mathcal{B}) \subseteq \{E_i\}$ for some $\{E_i\} \in \mathcal{A}$). We now introduce our continuity axiom.

A2: (Σ_{E_i} -continuity) For any $M \in \Sigma_{E_i}$ the sets $\{\phi_{M'} : M' \in \Sigma_{E_i}, M' \succeq M\}$ and $\{\phi_{M'} : M' \in \Sigma_{E_i}, M' \preceq M\}$ are (relatively) closed.¹⁴

Notice that this is indeed weaker than continuity since we only require continuity of the preferences along sequences of menus which are measurable with respect to the same (finite) Borel partition. For readers concerned that the notion of Σ_{E_i} is not definable from behavioral primitives, we make two remarks. First, we can easily define the notion of Σ_{E_i} -measurability using only \succeq . Fix a Borel partition and consider the following relation, \mathcal{R} , on the (objective) states. Say that $s \mathcal{R} s'$ iff $\sup(M(s)) \sim \sup(M(s'))$. Recall that $\sup(M)$ denotes the $\succeq |_{\Delta(X)}$ -maximal elements of M . The relation \mathcal{R} is clearly an equivalence relation on \mathcal{S} . Call a menu M

¹³Note that hyperplanes have non-zero intersection, so that we need a consistent way to assign boundary states to pairs of cells (E_i, E_j) that share a hyperplane on their boundary. This turns out to be an important issue that we revisit in the course of the proof of Theorem 2.

¹⁴The topology is the induced Euclidean topology on the space of utility vectors $\{\phi_M : M \in \Sigma_{E_i}\}$. We also note that this axiom is strictly weaker than the standard Hausdorff continuity axiom, listed as A2 in the list of DLR axioms. The proof that this is the case can be read off from our proof that A2 is necessitated by the Bewley model.

Σ_{E_i} -measurable if the partition induced by this equivalence relation is a coarsening of $\{E_i\}$.

Second, note that while the closed condition is stated on the set of value functions induced by menus, it is a simple exercise to check that A2 is equivalent to Hausdorff continuity on the set of $\{E_i\}$ -measurable menus. That is, the set $\{M' \in \Sigma_{E_i} : M' \succeq M, M \in \Sigma_{E_i}\}$ is closed (and similarly for the lower contour sets). Combining A0–A4 yields the following generalization of Bewley’s result to the DLR model. Let $\Sigma_{\mathcal{A}}$ denote the set of all \mathcal{A} -measurable menus (in $M(X)$).

Definition: A pair (u, Π) , where $u : \Delta(X) \rightarrow \mathbf{R}$ and $\Pi \subseteq \Delta(\mathcal{S})$, is called a *subjective Bewley representation* of a preference $(\succeq, \Sigma_{\mathcal{A}})$ if we have

$$M \succeq M' \Leftrightarrow E_{\pi} \phi_M \geq E_{\pi} \phi_{M'}, \forall \pi \in \Pi$$

The following is the main representation result of the paper.

Theorem 2 (Main Result). *A preference $(\succeq, \Sigma_{\mathcal{A}})$ admits a subjective Bewley representation (u, Π) if and only if it satisfies A0 – A4. Moreover, if (u', Π') is any other subjective Bewley representation, then $\Pi' = \Pi$ and $u' = \alpha \cdot u + \beta$ for some $\alpha > 0, \beta \in \mathbf{R}$.*

The following corollary provides a “subjective” Anscombe-Aumann representation for menu preferences which satisfy the weakened continuity axiom A2.

Corollary 1. *An order $(\succeq, \Sigma_{\mathcal{A}})$ satisfies A2–A4 if and only if it admits a subjective Bewley representation (u, Π) , where $\Pi = \{\pi\}$ is a singleton set.*

Note that this generalizes the classical Anscombe-Aumann theorem in two directions. First, the objective state space is not finite. Second, the support of the preference relation is the set of menu-induced AA acts on this state-space, not the full set of objective acts. Also note that we have only provided a representation for menus in $\Sigma_{\mathcal{A}}$, i.e. menus whose associated value functions ϕ_M have finite support.¹⁵ An extension of the representation to the full space $M(X)$ is currently work in progress. We make some (at this point, completely speculative) remarks on how a putative extension argument would proceed in the appendix. I now verify that A2 is implied by the subjective Bewley model.

Proposition 1 (Necessity of A2). *If \succeq admits a subjective Bewley representation, then it satisfies A2 (Σ_{E_i} -continuity).*

¹⁵For this reason, I will at times refer to the representation result (for each fixed space Σ_{E_i}) as the ‘finite subjective Bewley’ theorem.

Proof. The necessity argument proceeds in three steps. First, we use a closed graph argument to show that the limiting utility value of a sequence of menus evaluated at any state s is a lower bound on the value of the limiting menu at state s . Second, we use this argument to verify that the set Σ_{E_i} is closed (in \mathbf{R}^k , where $k = |\{E_i\}|$). Third, we use this to verify that A2 holds.

Step 1: A Closed Graph Argument.

Fix $M_n \in \Sigma_{E_i}$ and assume the sequence converges to some M^* . We check that $M^* \in \Sigma_{E_i}$ (the key part of the argument for this is in this step, but the formal argument spills into step 2 as well). Let $\{E_1, \dots, E_k\}$ be an enumeration of the partition cells and let $\{u_1^n, \dots, u_k^n\}$ be an enumeration (with possible repetition) of the values of ϕ_{M_n} evaluated at these cells. Passing to a common sub-sequence if necessary we can assume that the values u_i^n all converge. Now let $\Theta := \{1/n\}_{n=1}^\infty \cup \{0\}$ and consider the correspondence $D_s(1/n) := \arg \max_{f \in M_n} u_s(f)$. A straightforward application of Berge's Theorem of the Maximum reveals that this is an upper hemicontinuous correspondence $D_s : \Theta \rightrightarrows \Delta(X)$. Take the function $u_s(1/n, f) := u_s(f)$ and $D(1/n) := M_n, D(0) := M^*$. Note that $u_s(\cdot, \cdot)$ is continuous on $\Theta \times \Delta(X)$ and $D(\cdot)$ is a continuous (i.e. uhc and lhc) correspondence, so that the hypotheses of Berge's Theorem are met. The proof that $D(\cdot)$ is a continuous correspondence is straightforward, but since we have not been able to locate a reference we provide a proof here for completeness. Recall the following notation. Let $d_E(\cdot, \cdot)$ denote the Euclidean metric on $\Delta(X)$ and let $d_H(\cdot, \cdot)$ denote the Hausdorff metric on $M(X)$. Let $\Theta := \{1/n\}_{n=1}^\infty \cup \{0\}$ (viewed as a topological space with the relative topology from $[0, 1]$) and let $M_n \in M(X)$ be such that $d_H(M_n, M^*) \rightarrow 0$. Define a correspondence, $D : \Theta \rightrightarrows \Delta(X)$ via $D(x) = M_{1/x}$ (for $x \in \Theta$), where $D(0) := M^*$.

Claim 1. $D(\cdot, \cdot)$ is uhc and lhc (hence, continuous).

Proof of Claim 1. Since any $x (\neq 0) \in \Theta$ is isolated, it is easy to check that $D(x)$ is uhc and lhc. We verify that this also holds at $x = 0$. First we verify uhc. Let $V \subseteq \Delta(X)$ be an open set with $M^* \subseteq V$. For each $v \in V$ find an $\epsilon_v > 0$ so that $B_{2\epsilon_v}(v) \subseteq V$ (i.e. the open ball of radius $2\epsilon_v$ is contained in V). Note that the collection $\{B_{\epsilon_v}(v)\}_{v \in V}$ covers M^* . Since M^* is compact, extract a finite subcover, $B_{\epsilon_{v_1}}(v_1), \dots, B_{\epsilon_{v_n}}(v_n)$. Put $M' := \cup_{i=1}^n B_{\epsilon_{v_i}}(v_i)$ and find $N > 0$ large enough so that $d_H(M_n, M^*) < \min\{\epsilon_{v_i} : i = 1, \dots, n\}, \forall n \geq N$ and let $U := \{1/n : n \geq N\}$. Let $x \in U$ and take any $t \in M_{1/x}$. I claim that there is some $v_i \in \{v_1, \dots, v_n\}$ such that $d_E(v_i, t) < 2\epsilon_{v_i}$. To see this note that $d_H(M_{1/x}, M^*) < \epsilon_{v_i}, \forall i$ implies that there some $z \in M^*$ such that $d_E(t, z) < \epsilon_{v_i}, \forall i$. Find a v_i such that $z \in B_{\epsilon_{v_i}}(v_i)$ (since M' covers M^*) and note that

$$d_E(v_i, t) \leq d_E(v_i, z) + d_E(z, t) < \epsilon_{v_i} + \epsilon_{v_i}$$

Thus, $t \in B_{2\epsilon_{v_i}}(v_i)$. It follows that $M_{1/x} \subseteq \cup_{i=1}^n B_{2\epsilon_{v_i}}(v_i)$. Since this holds for any $x \in U$ and $B_{2\epsilon_{v_i}}(v_i) \subseteq V$, it follows that $D(x) \subseteq V, \forall x \in U$. This verifies uhc. For

lhc, let V be open with $V \cap M^* \neq \emptyset$ and find $\epsilon > 0$ with $v \in M^* \cap V$ and $B_\epsilon(v) \subseteq V$. Find $N > 0$ such that $d_H(M_n, M^*) < \epsilon, \forall n \geq N$ and put $U := \{1/n : n \geq N\}$. Since $d_H(v, M_n) \leq d_H(M^*, M_n) < \epsilon, \forall n \geq N$ there is some $t_n \in M_n$ with $d_E(t_n, v) < \epsilon$. Thus, $t_n \in V$ - so that $M_n (= D(1/n)) \cap V \neq \emptyset, \forall n \geq N$. \square

Return to the proof of the Proposition. Now consider the graph of the correspondence, $D_s(\cdot)$, in $\Theta \times \Delta(X)$, where we set

$$\text{Gr}(D_s) := \{(x, y) \in \Theta \times \Delta(X) : y \in D_s(x)\}$$

Take $s \in E_i$ and consider a sequence of pairs $(1/n, f_n)$ where $f_n \in D_s(1/n)$ is such that $u(f_n) = u_i^n$. Passing to a convergent subsequence if necessary we obtain a convergent sequence $(1/n, f_n) \in \Theta \times \Delta(X)$. By the Closed Graph Theorem, the limit of this sequence $(0, f^*)$ is an element of $\text{Gr}(D_s)$.¹⁶ Note that $\lim u_i^n = \lim u(f_n) = u(\lim f_n) = u(f^*)$. Thus, setting $u_i^* := u(f^*)$ we obtain the following conclusion:

$$\text{For every } s \in E_i, \exists f \in \arg \max_{g \in M^*} u_s(g) \text{ s.t. } u(f) = u_i^*.$$

Step 2: Σ_{E_i} is closed.

Now consider a state $s \in E_i$ and, abusing notation, let s_u denote the state with cardinal utility $u(\cdot)$. Let $M_s^* = \arg \max_{f \in M^*} u_s(f)$ and let $\bar{u} := \max_{f \in M_s^*} u(f)$. The preceding observation implies that $\bar{u} \geq u_i^*$. We check the reverse inequality. Consider first the case where s is an interior state. Find $\alpha \in (0, 1)$ so that $s_\alpha := \alpha \cdot s + (1 - \alpha) \cdot s_u \in E_i$. Let $M_{s_\alpha}^* := \arg \max_{f \in M^*} u_{s_\alpha}(f)$. Observe that, by definition of s_α , we must have $u(f) \geq \bar{u}, \forall f \in M_{s_\alpha}^*$.¹⁷ Applying the preceding closed graph argument to the state s_α we know that there is some $f \in M_{s_\alpha}^*$ with $u(f) = u_i^*$, so that $u_i^* \geq \bar{u}$. It follows that $\bar{u} = u_i^*$. Now consider the case where $s \in E_i$ is a border state. If there is an α with $s_\alpha \in E_i$, then we can apply the above argument verbatim. Otherwise, $s_\alpha \in E_j$ where the cell E_j is adjacent to E_i . Since the border of the pair (E_i, E_j) is assigned to the cell E_i , it must be the case that $u_i^* \geq u_j^*$. Take $\bar{u} = \max_{f \in M_s} u(f)$ as above and note that the closed graph argument yields the inequality $\bar{u} \geq u_i^*$. For the reverse, the preceding argument yields $u(f) \geq \bar{u}, \forall f \in M_{s_\alpha}^*$. By the closed graph argument applied to the state s_α , there is some $f \in M_{s_\alpha}^*$ such that $u(f) = u_j^*$. Thus, $u_j^* = u(f) \geq \bar{u}$. Since $u_i^* \geq u_j^*$ it follows that $u_i^* \geq \bar{u}$, implying the equality $u_i^* = \bar{u}$ when s is a border state. It follows that $M^* \in \Sigma_{E_i}$. Now note that elements of Σ_{E_i} are really equivalence classes of menus (all of which are equivalent under \succeq). Thus, closedness of Σ_{E_i} means that if we have a convergent sequence of equivalence classes $[M_n]$, then the limiting class $[M^*]$ is in Σ_{E_i} .¹⁸ To distinguish sequences of classes of menus from sequences of menus, let ϕ_M^n denote a

¹⁶Since D_s is uhc and closed-valued (and $\Delta(X)$ is a metric space, hence trivially Hausdorff) the hypotheses of the Closed Graph Theorem are satisfied.

¹⁷Else, if $f \in M_{s_\alpha}^*$ with $u(f) < \bar{u}$ find $f_{\bar{u}} \in M_s^*$ with $u(f) = \bar{u}$ and notice that $s_\alpha \cdot f = \alpha s \cdot f + (1 - \alpha) s_u \cdot f < \alpha s \cdot f_{\bar{u}} + (1 - \alpha) s_u \cdot f_{\bar{u}}$.

¹⁸The notation $[M]$ denotes the Σ_{E_i} class of the menu M .

convergent sequence in Σ_{E_i} - with limit ϕ_M^* . Extract from this sequence of functions a sequence of menus, M_n , which are representatives of the classes corresponding to the functions $\phi_{M_n}^n(\cdot)$. Since the space $M(X)$ is a compact metric space under the Hausdorff metric, pass to a convergent subsequence $M_{n_k} \rightarrow M'$. By the preceding argument, $\phi_{M_{n_k}} \rightarrow \phi_{M'} \in \Sigma_{E_i}$. Moreover, $\phi_{M_{n_k}} = \phi_{M'}^{n_k} \rightarrow \phi_M^*$ since the sequence $\phi_{M'}^n$ is convergent. Thus, $\phi_M^* = \phi_{M'}$ - showing that Σ_{E_i} is closed.

Step 3: Verification.

To verify axiom A2 let $M_n \rightarrow M$ where $M_n \succeq M'$ and $M_n, M' \in \Sigma_{E_i}$. By step 1, $M \in \Sigma_{E_i}$. Let $\{u_1^n, \dots, u_k^n\}$ be the values assumed by the function $\phi_{M_n}(\cdot)$ on the cells E_1, \dots, E_k and let $\{u_{M'}(i)\}$ denote this set of values for $\phi_{M'}$. Note that representation by the Bewley model means: (abusing notation we revert to letting Θ denote the set of measures in the Bewley representation)

$$M_n \succeq M' \Leftrightarrow \sum_{i=1}^n u_i^n \cdot \pi_i \geq \sum_{i=1}^n u_{M'}(i) \cdot \pi_i, \forall \pi \in \Theta$$

Now choose *any* subsequence M_{n_j} such that the associated k -tuple of utility values $\{u_1^{n_j}, \dots, u_k^{n_j}\}$ converges.¹⁹ It follows that

$$\sum_{i=1}^k u_i^{n_j} \cdot \pi_i \geq \sum_{i=1}^k u_{M'}(i) \cdot \pi_i \quad \forall n_j \Rightarrow \sum_{i=1}^k u_i^* \cdot \pi_i \geq \sum_{i=1}^k u_{M'}(i) \cdot \pi_i$$

Since this holds for all $\pi \in \Theta$ (and we are assuming \succeq admits the above Bewley representation), it follows that $M \succeq M'$. \square

To place our continuity axiom in context let us recall the Strotz model, axiomatized in Gul-Pesendorfer GP (2005). This model is comprised of two maps, $u, v : \Delta(X) \rightarrow \mathbf{R}$ which are assembled into a menu utility as follows:

$$U^{ST}(A) := \max_{x \in A_v} u(x)$$

where $A_v := \arg \max_{z \in A} v(z)$. It is well-known that the Strotz model does not satisfy Hausdorff continuity. Consider the following example, taken from GP (2001). To fix ideas, think of $x = \text{work}$, $y = \text{shirk}$, $z = \text{play}$. The DM wants to commit to work but is tempted to play. Model this with a Strotz utility as follows. [Utility Drop Problem] Let $u(x) > u(y) > u(z)$, $v(z) > v(x) = v(y)$ and make the obvious extension to lotteries on $\{x, y, z\}$. Put $M = \{x, y\}$, $y_\epsilon = (1 - \epsilon) \cdot y + \epsilon \cdot z$, and $M_\epsilon = \{x, y_\epsilon\}$. Note that $U^{ST}(M) = u(x)$ and $U^{ST}(M_\epsilon) = u(y_\epsilon)$ so that $U^{ST}(M)$ is bounded above $U^{ST}(M_\epsilon)$ even as $M_\epsilon \rightarrow M$. Since the subjective Bewley nests the Strotz model,

¹⁹Note that the argument in step 1 yields that the limiting utility values u_i^* are the same no matter which sub-sequence we pick - so that the limit function ϕ_M is independent of the choice of subsequence.

it is subject to the same ‘utility drop problem’ described in the example. How does Σ_{E_i} -continuity get around this problem? The idea is that it only requires continuity along sequences M_n (with $M_n \rightarrow M$) where every element of the sequence is measurable with respect to the *same* Borel partition. To see this concretely, let us check that the example given above is consistent with A2. I claim that no sequence of menus M_{ϵ_n} can belong to the same Borel partition. Check this via contradiction: Let $\{E_i\}$ denote a doubleton partition such that $M_\epsilon \in \Sigma_{E_i}$ for infinitely many M_ϵ . Note that, for this doubleton partition, we may group states into three sets, $\Sigma_1(\epsilon) := \{s \in \mathcal{S} : y_\epsilon \succ_s x\}$, $\Sigma_2(\epsilon) := \{s \in \mathcal{S} : x \succ_s y_\epsilon\}$, $\Sigma_3 := \{s \in \mathcal{S} : x \sim_s y_\epsilon\}$. Note that if M_{ϵ_n} is Σ_{E_i} -measurable for each ϵ_n , then the sets $\Sigma_1(\epsilon_n), \Sigma_2(\epsilon_n), \Sigma_3(\epsilon_n)$ are independent of n . Now consider a state, s_v , that corresponds to the ranking v in the UDP example, so that $y_\epsilon \succ_{s_v} x$ for all ϵ . Then, $s_v \in \Sigma_1(\epsilon_n), \forall n$.

OTOH, the argument in step 1 of the proof of Proposition 2 shows that if $M_{\epsilon_n} \rightarrow M$ in the Hausdorff metric and $M_{\epsilon_n} \in \Sigma_{E_i}$, then $M \in \Sigma_{E_i}$. Now notice that since $x \neq y$, we can find some state $s \in \mathcal{S}$ such that $x \not\sim_s y$.²⁰ Thus, M induces the same doubleton partition $\{E_i\}$ of the state space as the menus M_{ϵ_n} . But this gives a contradiction. The state s_v must lie on the hyperplane that defines the partition since $x \sim_{s_v} y$. Yet, for each of the menus M_{ϵ_n} we have $y_{\epsilon_n} \succ_{s_v} x$ - so that the hyperplanes defining the partitions generated by (resp.) M and M_{ϵ_n} must be distinct.

2.1 A Subjective States Axiomatization

Notice that the state space \mathcal{S} is explicitly referenced in our axioms - namely, A2 and A4. This breaks with the subjective state space approach of most menu choice papers. Our motivation for this model is to see if richness (resp. sparseness) of the choice domain presents a *source* of incompleteness. In particular, we have in mind the following comparative statics exercise. Imagine we have two decision-makers, DM1 and DM2. The support of the preference relation of the first DM is the full domain of Anscombe-Aumann acts. The support of the preference relation of the second DM is the subset of this domain consisting of Anscombe-Aumann acts induced by menus. Moreover, the preference relation of the first DM is assumed to

²⁰To see this note that by the Separating Hyperplane Theorem, there is a vector $p \in \mathbf{R}^k$ such that $p \cdot x > c > p \cdot y$. Since x, y are lotteries, putting $\vec{c} = c \cdot \vec{1}$ we obtain, $(p - \vec{c}) \cdot x > (p - \vec{c}) \cdot y$. Now find a constant K big enough so that $(p - \vec{c})_i + K > 0$, i.e. every component of the vector $p - \vec{c} + \vec{K}$ is positive. Note that $(p - \vec{c} + \vec{K}) \cdot x = (p - \vec{c}) \cdot x + K > (p - \vec{c}) \cdot y + K = (p - \vec{c} + \vec{K}) \cdot y$. Divide through by $\sum_i (p - \vec{c} + \vec{K})_i$ and put

$$s_* := \frac{(p - \vec{c} + \vec{K})_i}{\sum_i (p - \vec{c} + \vec{K})_i}$$

We then have $x \succ_{s_*} y$.

be an extension of the preference relation of the second DM. If indeed the ‘size’ of the choice domain provides a source of incompleteness - in the sense that the second DM has a more incomplete relation - and both preferences are represented by the Bewley model, then the collection of measures that captures incompleteness for DM2 is a superset of DM1’s measures. The point of the representation result is that it proves the Bewley representation for DM2, so that we can use the same decision model - and a common parameter - to measure the relative incompleteness of the decision-makers’ preference relations.

This is why we have expressed the axioms of the model with an explicit (i.e. objective) state space. Nevertheless, it is still legitimate to ask - can all this be done with a subjective state space? In other words, imagine DM2’s relation on menus is still represented by the Bewley model and DM1’s relation is still an extension of this model. However, the state space that DM2 visualizes as the relevant space of uncertainty is not outwardly observable to the modeler. Thus, we seek axioms on menus that are equivalent to the given A0 – A4, but that make no explicit mention of a state space. To carry this out, define the following axioms. Let $\inf(A)$ (resp. $\sup(A)$) denote the infima (resp. suprema) w.r.t to $\succeq|_X$ (the singleton ranking) on the (closed) menu A .

A4*: (Weak Dominance) If $\inf(A) \succeq \sup(B)$, then $A \succeq B$.

Let us interpret. In the unobserved second stage of choice, some item from the menu is chosen. Moreover, interpreting via the Bewley model, states are unobserved, but state-dependent preferences over lotteries are the same across states. Hence, if menu A has the property that its worst lottery (where ‘worst’ is determined by the state-independent utility) is weakly better than the best lottery in B , then the DM should prefer A to B .²¹ Now we reformulate the partition-dependent Hausdorff continuity axiom without explicitly mentioning the state space. For the reader who has seen A2 already, this construction will seem a little contrived. We agree - it is a contrivance. In our view, the importance of the translation of the axioms is mostly symbolic - it is just to show that our primitive is really no different than the standard menu choice primitive (despite the fact that two of our axioms have states explicitly in the statement).

Given a menu M , let $\{I_u(f)\}_{f \in M}$ denote the partition of M , where elements of each partition cell are $\succeq|_X$ -indifference classes (w.r.t the ranking on singleton menus applied to M). Let $\{I_u^1, \dots, I_u^k\} = M$ be an enumeration of this partition

²¹Note that in the objective Anscombe-Aumann setting (and with a finite state space), Weak Dominance can replace the usual monotonicity axiom and, along with the other Bewley axioms (and Indifference - see A5 below), still characterizes the Bewley model.

(ranking indifference classes from top to bottom). For any fixed lottery $p \in \Delta(X)$, we can compute the dot product $p \cdot f$ for each $f \in M$. Say that $p \in I_u(i)$ (not to be confused with I_u^i) if $p \cdot f$ is maximal for some $f \in I_u^i$. If this happens for more than one partition cell, break ties in the direction of the singleton ranking, i.e. choose the lower-numbered indifference class. Note that the partition, $\{I_u^1, \dots, I_u^k\}$, of M now induces a partition, $\{I_u(1), \dots, I_u(k)\}$, of $\Delta(X)$, call it Π_M . Let HC = Hausdorff continuity and let $\Sigma_\Pi := \{M \in M(X) : \Pi_M = \Pi = \Pi_{M'}\}$.

A2*: (Partition-Dependent HC) Fix any Σ_Π . Then, for any $M \in \Sigma_\Pi$ the relative upper-contour (resp. lower-contour) set

$$\{M' \in M(X) : M' \succeq M\} \cap \Sigma_\Pi$$

is closed in Σ_Π .

Finally, we require one more axiom for a subjective state space axiomatization.

A5: (Indifference) If $M, M' \in \Sigma_\Pi$ have the property that $\{f\} \sim \{f'\}$ for any $f \in M \cap I_u^i, f' \in M' \cap I_u^i$, then $M \sim M'$.

In words, take two menus which induce the same partition Σ_Π of lottery space. If the second stage of choice yields the same utility from both menus, for every possible second stage choice, then the DM is indifferent between the two menus. The observation below just records that $A0, A1, A2^*, A3, A4^*, A5$ are equivalent to $A0 - A4$. The proof of the equivalence goes via the proof of the representation theorem. Clearly, if we think of a menu as an implicit Strotzian value function, then $A2^*$ itself is equivalent to $A2$. The less trivial observation is that we can switch our $A4$ for $A4^*$ (weak dominance) and still obtain the same representation. To show this, we go via the sufficiency argument presented below. There is a step in the proof of the direction “ $A0 - A4 \Rightarrow$ Bewley representation” where we will invoke $A4$ (namely, [here](#)). In this step, we will point out that (in the presence of the other axioms) replacing $A4$ with $A4^*$ yields the same implication. The proof transfers the menu preference to an induced relation on the space of Strotzian value functions. Indifference ensures that this transfer is well-defined.

Observation 1. *A partial order (\succeq, Σ_A) satisfies $A0 - A4$ if and only if it satisfies $A0, A1, A2^*, A3, A4^*, A5$.*

3 Proof of Theorem 2

Recall that the space \mathcal{S} is a superset of the set, Π_{vNM} , of vNM preferences on $\Delta(X)$. Let $\Pi^* : \mathcal{S} \rightarrow \Pi_{\text{vNM}}$ be the natural projection map defined by $\Pi^*(s) = \succeq_s$, that

is, the projection maps a state s to the ordinal vNM preference generated by that state. Put $[s] = (\Pi^*)^{-1}(s)$. We say that a Borel partition, $\{E_i\}$, is *coherent* if for any $s \in \mathcal{S}$, whenever $s \in E_i$, then $[s] \subseteq E_i$. In the course of the proof of Theorem 2, we will define sets of measures for each Borel partition of the state space. As we refine the partitions, we will see that the sets of measures are also ‘refined’ in a corresponding fashion. Thus, for a fixed partition of the state space we inductively define a sequence of refinements with an associated sequence of refined measures on these partitions. The importance of the Borel property is that it allows us to show that the limiting process, particularly the set of measures obtained in the limit, are *independent* of the refinement process. On the other hand, we would also like to only consider coherent Borel partitions since these are the only partitions that allow us to interpret the results in terms of the set of ex-post vNM preferences. Note that a Borel partition is specified by a collection of pairs $\{(p_i, c_i)\}$ consisting of a separating vector together with a separating constant. By choice of the space \mathcal{S} we can equivalently describe the partition with a collection $\{(p_i^*, c_i^*)\}$, where $p_i^* \in \Delta(X)$. That is, we can take the separating vectors to be lotteries. We refer to the collection $\{p_i^*, c_i^*\}$ as a *lottery representation* of the given Borel partition $\{E_i\}$. The following proposition gives a complete, concrete characterization of the set of Borel partitions in terms of their respective lottery representations.

Proposition 2. *Let $\{E_i\}$ be a Borel partition of \mathcal{S} . The following are equivalent,*

1. $\{E_i\}$ admits a lottery representation $\{p_i^*, c_i^*\}$, where $c_i^* \equiv \frac{1}{k}, \forall i$.
2. $\{E_i\}$ is coherent.

Proof. (2) \Rightarrow (1): Let $\{p_i, c_i\}_{i \in I}$ be a Borel partition that generates $\{E_i\}$ and note that $\{E_i\}$ is coherent if and only if the doubleton partition, $\mathcal{S}_1(i) := \{s : p_1 \cdot s \geq c\}$, $\mathcal{S}_2(i) = \mathcal{S} \setminus \mathcal{S}_1(i)$, is coherent for each $i \in I$. Thus, we are reduced to the case of doubleton partitions, say $\mathcal{S}_1 = \{s : p \cdot s \geq c\}$, $\mathcal{S}_2 = \{s : p \cdot s < c\}$. Note that we may normalize p such that $p_i \geq 0$ and $\sum_i p_i = 1$, so that lottery representations always exist. Via contraposition, if we have a lottery representation of the partition with $c \neq 1/k$, say $c > 1/k$. then consider

$$s' := \alpha \cdot s + (1 - \alpha) \cdot \vec{\frac{1}{k}}$$

where $s \in \mathcal{S}_1$ and $\vec{\frac{1}{k}} = (1/k, 1/k, \dots, 1/k)$. Note that by coherence, $s' \in \mathcal{S}$. Choose α sufficiently close to 0 so that $p \cdot s' < 1/k + \epsilon < c$ - contradiction.

(1) \Rightarrow (2): As above we reduce to doubleton partitions $\mathcal{S}_1 := \{s : p \cdot s \geq c\}$, $\mathcal{S}_2 = \{s : p \cdot s < c\}$, where p is a lottery and $c = 1/k$ by hypothesis. Define the following *companion* lottery, $p' := \frac{(\vec{1}-p)}{k-1}$ and consider the following coherent partition of \mathcal{S} ,

$$E_1^* := \{s : p \cdot s \geq p' \cdot s\}, E_2^* := \{s : p \cdot s < p' \cdot s\}$$

We claim that $\mathcal{S}_1 = E_1^*, \mathcal{S}_2 = E_2^*$. To see this note that

$$p \cdot s \geq p \cdot s' = \frac{(\vec{1} - p)}{k - 1} \cdot s = \frac{1}{k - 1} - \frac{p \cdot s}{k - 1}$$

so that $s \in E_1^*$ iff $p \cdot s \geq 1/k$ iff $s \in \mathcal{S}_1$. \square

Transfer to Ex Post Utility Space

We transfer menu preferences to the space of induced ex post utility vectors, i.e. Strotzian value functions. Let $\{E_i\}_i$ denote a (finite) Borel partition of \mathcal{S} and, abusing notation, let $\{E_i\}$ also denote the σ -algebra on \mathcal{S} generated by this partition. Let Σ_{E_i} denote the set of $\{E_i\}$ -measurable menus, i.e. consider the function $\phi_M : \mathcal{S} \rightarrow \mathbf{R}$ defined by $s \mapsto \max_{f \in M_s} u(f)$, where $M_s := \arg \max_{f \in M} u_s(f)$. Consider the σ -algebra defined by $\sigma(\phi_M^{-1}(B) : B \in \mathcal{B})$, where \mathcal{B} is the Borel sigma-algebra on \mathbf{R} . The menu M is said to be $\{E_i\}$ -measurable if $\sigma(\phi_M^{-1}(B) : B \in \mathcal{B}) \subseteq \{E_i\}$. Let $k = |\{E_i\}|$ be the cardinality of the partition $\{E_i\}$ and identify the set Σ_{E_i} with the equivalence classes of vectors in \mathbf{R}^k under the map into \mathbf{R}^k implicitly defined by $M \mapsto \phi_M$.²²

Denote the image in \mathbf{R}^k of the map $M \mapsto \phi_M$ by \mathcal{U}_ϕ . Notice that \mathcal{U}_ϕ is a closed and convex set, and that it inherits from Σ_{E_i} (as a consequence of A4 and the fact that the map $M \mapsto \phi_M$ is affine, i.e. $\phi_{\alpha \cdot M + (1-\alpha) \cdot M'} = \alpha \cdot \phi_M + (1-\alpha) \cdot \phi_{M'}$) the pre-order on menus that satisfies axioms A0 – A4 (equivalently, A0, A1, A2*, A3, A4*). I claim that $\dim(\mathcal{U}_\phi) = k$, the set \mathcal{U}_ϕ contains a basis for the vector space \mathbf{R}^k .

Aside: Monotonicity v.s. Weak Dominance

The way we check the dimension claims is by producing a collection of menu-induced vectors ϕ_M^i that generate a basis. In particular, starting with a fixed vector ϕ_M we produce (by induction) perturbed vectors ϕ_M^i where ϕ_M^i is equal to ϕ_M in all coordinates $j \neq i$ and strictly bigger in coordinate i . By A4, this means the difference $\Phi_M^i - \Phi_M$ is in the pointed cone at the origin. So that the differences $\{\Phi_M^i - \Phi_M\}$ span the full vector space. Now assume we only assumed A4* - does this same step carry through? To see that it does, let \underline{M} denote a menu consisting of a singleton $\succeq_{|X}$ -minimal lottery and note that we have, by Weak Dominance, $\phi_M^i - \phi_{\underline{M}}$ in the pointed cone. Hence, $\Phi_M^i - \phi_{\underline{M}} - (\phi_M - \phi_{\underline{M}}) = \phi_M^i - \phi_M$ is in the span of the pointed cone. Since the span has full dimension, this is all that is needed for the next step.

²²In defining this map, there is a subtle point concerning the borders of the partitions $\{E_i\}$. However, this issue can be skirted by noting that a menu can be identified with the vector ϕ_M where we just consider the value of the function ϕ_M at interior states. It is a simple exercise to check that this association is well-defined. Thus, we can think of the set \mathcal{U}_ϕ as the collection of vectors ϕ_M , where we only evaluate the functions ϕ_M at interior states.

Now assume we have shown that $\dim(\mathcal{U}_\phi) = k$. By convexity, this would then imply that \mathcal{U}_ϕ has non-empty interior. Think of the coordinates on the vectors in \mathcal{U}_ϕ as a space of symbols, $\{1, 2, \dots, n\}$. Standard arguments then imply the existence of a set of probability measures $\Theta \subseteq \Delta(\{1, 2, \dots, n\})$ such that for any two vectors $(u_M, u_{M'}) \in \mathcal{U}_\phi$ (with $u_M := (u_M(1), \dots, u_M(n))$ and $u_{M'} := (u_{M'}(1), \dots, u_{M'}(n))$) we have

$$u_M \succeq u_{M'} \Leftrightarrow \sum_{i=1}^n u_M(i) \cdot \pi_i \geq \sum_{i=1}^n u_{M'}(i) \cdot \pi_i, \forall \pi \in \Theta$$

Now pull back the measures in Θ to the state space \mathcal{S} via the map $M \mapsto \phi_M$ to obtain (abusing notation, refer to a generic pullback as π as well)

$$M \succeq M' \Leftrightarrow \sum_{i=1}^n u_M(i) \cdot \pi(s : \phi_M(s) = u_M(i)) \geq \sum_{i=1}^n u_{M'}(i) \cdot \pi(s : \phi_{M'}(s) = u_{M'}(i)), \forall \pi \in \Theta$$

There is a subtle point involved in making sure the above pull-back construction is well-defined. First, note that in setting $\pi(s : \phi_M(s) = u_M(i)) = \pi(i)$ we are implicitly assuming that the *only* states where the value of ϕ_M equals $u_M(i)$ are those states that comprise the partition cell E_i . This means that if s is any state that lies at the border of two cells, say, (E_i, E_j) , then if $M \in \Sigma_{E_i}$ is such that $u_M(i) = \phi_M(s) \geq u_M(j)$, then for *any* other menu $M' \in \Sigma_{E_i}$ we must also have $u_{M'}(i) = \phi_{M'}(s) \geq u_{M'}(j)$. Thus, consider the collection of interiors of the partition cells, $\text{int}(E_i)$. The choice of a partition $\{\hat{E}_i\}$ whose interiors correspond the collection $\{\text{int}(E_i)\}$ implicitly determines an assignment rule on boundary states. The preceding sentence requires that we are able to find a partition with an assignment rule satisfying the consistency condition: (denote the border of a cell E_j as ∂E_j)

$$\text{If } s \in E_i \text{ and } s \in \partial E_i \cap \partial E_j, \text{ then } u_M(i) = \phi_M(s), \forall M \in \Sigma_{E_i}.$$

To show that there is (in fact, canonical!) an assignment rule satisfying the above consistency condition I break the forthcoming argument into two pieces. First, I reduce to the case where the collection of hyperplanes $\{\mathcal{L}_1, \dots, \mathcal{L}_k\}$ which generate the partition all have the property that $s_u \notin \mathcal{L}_i$, i.e. the state corresponding to the cardinal ranking on singletons, $u(\cdot)$, is in the interior of the partition. Second, I verify that partitions formed by hyperplanes which contain s_u are observable indistinguishable from those obtained by deleting these hyperplanes. For this reason, I refer to the steps (resp.) as (i) reduction and (ii) removal of unobservable hyperplanes.

Reduction

Consider the following assignment rule. Abusing notation, let $\{\text{int}(E_i)\}_i$ denote the (disjoint) collection of sets formed by the intersections of the hyperplanes $\mathcal{L}_1, \dots, \mathcal{L}_k$. We now specify a partition of \mathcal{S} where the set of interiors of the partition cells is precisely the collection $\{\text{int}(E_i)\}$. We assign boundary states as follows. For each state

s , let $\Sigma(s) = \{\mathcal{L}_i : s \in \mathcal{L}_i\}$. Consider the set of states $\{s'_\alpha : s' = \alpha s + (1 - \alpha)s_u\}$ (recall that s_u was the state whose vNM preference agrees with the normative ranking $u(\cdot)$). Notice that (since $s_u \notin \mathcal{L}_i$) each s'_α is in $\text{int}(E_i)$ for some i . Assign the state s to the unique $\text{int}(E_i)$ such that $s'_\alpha \in \text{int}(E_i), \forall \alpha \uparrow 1$. Note that if $s \in \text{int}(E_i)$, then s gets assigned to $\text{int}(E_i)$ under this rule. This rule yields an assignment function,

$$\kappa : \mathcal{S} \rightarrow \{\text{int}(E_1), \text{int}(E_2), \dots, \text{int}(E_n)\}$$

and put

$$E_i := \{s \in \mathcal{S} : \kappa(s) = \text{int}(E_i)\}$$

This defines a Borel partition. I now verify that it has the desired consistency property.

To check this take a pair of adjacent cells (E_i, E_j) (i.e. the cells share at least one hyperplane boundary) and let $M, M' \in \Sigma_{E_i}$. I claim that:

$$u_M(i) \geq u_M(j) \Leftrightarrow u_{M'}(i) \geq u_{M'}(j)$$

Find a state $s \in \partial E_i \cap \partial E_j$ such that there is an $\epsilon > 0$ with $B_\epsilon(s) \subseteq E_i \cup E_j$, i.e. the ϵ -ball in \mathcal{S} around state s lies entirely in the union of the two cells. Find a pair of sequences s_i^n, s_j^n such that $s_i^n \in E_i, s_j^n \in E_j$ and $s_i^n, s_j^n \rightarrow s$. Also find an associated pair of sequences of lotteries (f_i^n, f_j^n) such that $f_i^n \in \arg \max_{f \in M} u_{s_i^n}(f), f_j^n \in \arg \max_{f \in M} u_{s_j^n}(f)$ and $u(f_i^n) = u_M(i), u(f_j^n) = u_M(j)$. Put $\Theta(i) := \{s_i^n\}_n \cup \{s\}$ and consider the correspondence $\Phi(i) : \Theta(i) \rightrightarrows \Delta(X)$ given by $\Phi(i)(s) = \arg \max_{f \in M} u_s(f)$. Consider the graph of this correspondence, $\text{Gr}(\Phi(i)) = \{(x, y) \in \Theta(i) \times \Delta(X) : y \in \Phi(i)(x)\}$ and (as in the proof of Proposition 1), after passing to a convergent subsequence if necessary, we have $(s, f^*) \in \text{Gr}(\Phi(i))$ (where $f^* := \lim f_i^n$). Note that $u(f^*) = u_M(i)$. Now repeat the argument with the sequence s_j^n to obtain $(s, f_*) \in \text{Gr}(\Phi(j))$ where $f_* := \lim f_j^n$ and $u(f_*) = u_M(j)$. If $\kappa(s) \in E_i$ (i.e. state s is assigned to cell E_i), then note that we must have $u_M(i) \geq u_M(j)$. This is a similar argument as given in step 3 of Proposition 1, but we repeat for the reader's convenience. Find $\alpha \in (0, 1)$ such that $\alpha \cdot s_u + (1 - \alpha) \cdot s \in \text{int}(E_i)$. Towards contradiction, if $u_M(i) < u_M(j)$ then take an alleged $f_{s_\alpha} \in \arg \max_{f \in M} u_{s_\alpha}(\cdot)$ with $u(f_{s_\alpha}) = u_M(i)$. Note that

$$\begin{aligned} s_\alpha \cdot f_* &= \alpha s_u \cdot f_* + (1 - \alpha) s \cdot f_* \\ &> \alpha s_u \cdot f_{s_\alpha} + (1 - \alpha) s \cdot f_* \\ &\geq \alpha s_u \cdot f_{s_\alpha} + (1 - \alpha) s \cdot f_{s_\alpha} \quad (\text{since } f_* \in \arg \max u_s(\cdot)) \end{aligned}$$

This contradicts the assumption that $f_{s_\alpha} \in \arg \max u_{s_\alpha}(\cdot)$. Thus, $u_M(i) \geq u_M(j)$ whenever $\kappa(s) \in E_i$. A symmetric argument using f^* shows that $u_M(i) \leq u_M(j)$ whenever $\kappa(s) \in E_j$. Since this conclusion holds for all menus in Σ_{E_i} , the desired

consistency property follows.

Removal of Unobservable Hyperplanes

The preceding argument presumed that s_u was not on any of the hyperplanes, \mathcal{L}_i , that defined the partition. We verify that there are *no* observable implications to this assumption. Though the idea is simple, to formalize this ‘no observable implications’ condition requires some notation. From the preceding arguments we find that a (coherent) Borel partition can equivalently be described by a pair $(\{\mathcal{L}_i\}, \phi)$ consisting of a finite collection of hyperplanes, e.g. $\{s : p \cdot s = 1/k\}$, and an assignment rule $\gamma(\cdot)$ that assigns border states to cells. Although γ need not be consistent in the sense defined above, the state s must be assigned by $\gamma(s)$ to a cell for which it is part of the border. That is, let $\{\text{int}(E_{i_1}), \dots, \text{int}(E_{i_j})\}$ be a list of cells for which $s \in \overline{\text{int}(E_{i_j})}$. Then, we require that $\gamma(s) \in E_{i_j}$ for one of these j .

Now label the set of hyperplanes as $\{\mathcal{L}_1, \dots, \mathcal{L}_k, \mathcal{L}_{k+1}, \dots, \mathcal{L}_n\}$, where $s_u \notin \mathcal{L}_i$, $1 \leq i \leq k$ and $s_u \in \mathcal{L}_i$, $i \geq k+1$. Let $(\{\mathcal{L}_i\}_{i=1}^k, \kappa(\cdot))$ denote the Borel partition with assignment rule κ as defined above. Also note that the rule γ on the finer partition defined by the set $\{\mathcal{L}_1, \dots, \mathcal{L}_n\}$ induces an assignment rule on the (implicit) coarsening defined by $\{\mathcal{L}_1, \dots, \mathcal{L}_k\}$. Thus, starting with the consistent rule $\kappa(\cdot)$ defined on the set of planes $\{\mathcal{L}_1, \dots, \mathcal{L}_k\}$, find an extension of this rule γ that gives an assignment rule on the full set of hyperplanes, $\{\mathcal{L}_1, \dots, \mathcal{L}_n\}$ (subject to the constraint specified above). Take $\{E_i\}$ to be the Borel partition generated by the pair $(\{\mathcal{L}_i\}_{i=1}^n, \gamma)$ and let $\{E'_i\}$ be the coarser partition generated by the pair $(\{\mathcal{L}_i\}_{i=1}^k, \kappa)$. Let $M \in \Sigma_{E_i}$. I claim that, in fact, $M \in \Sigma_{E'_i}$. This is the formalization of the statement that there are no observable restrictions to considering only Borel partitions where all planes omit the state s_u . To check this, proceed by induction on the number of hyperplanes in the set $\{\mathcal{L}_{k+1}, \dots, \mathcal{L}_n\}$ - call this the set of unobservable hyperplanes.

Base Step: $|\{\mathcal{L}_{k+1}, \dots, \mathcal{L}_n\}| = 1$.

Let $\{E_i\}$ denote the Borel partition formed by all $k+1$ hyperplanes and let $\{E'_i\}$ denote the partition formed by $\{\mathcal{L}_1, \dots, \mathcal{L}_k\}$. Let $E'_{i_1}, \dots, E'_{i_j}$ be an enumeration of the partition cells that intersect the plane \mathcal{L}_{k+1} . Also let $E_{i_m}^1, E_{i_m}^2$ denote the two cells into which the plane \mathcal{L}_{k+1} splits E_{i_m} . Let $s \in E_{i_m}$ and as above find a pair of sequences $(s_n^1, s_n^2) \in E_{i_m}^1 \times E_{i_m}^2$ such that $s_n^1, s_n^2 \rightarrow s$. The same argument as above shows that $\phi_M(s) \geq \phi_M(s_n^1)$ and $\phi_M(s) \geq \phi_M(s_n^2)$. Say that $\phi_M(s_n^1) \geq \phi_M(s_n^2)$. Note that for all α close to 1 we have $s_\alpha := \alpha s + (1-\alpha)s_u \in \partial E_{i_m}^1 \cap \partial E_{i_m}^2$. Moreover, note that we have $u(f) \geq \phi_M(s), \forall f \in \arg \max u_{s_\alpha}(\cdot)$. Now apply the preceding closed graph argument again with $\hat{s}_n^2 \in E_{i_m}^2, \hat{s}_n^2 \rightarrow s_\alpha$ to find an $f^* \in \arg \max u_{s_\alpha}(\cdot)$ with $u(f^*) = \phi_M(\hat{s}_n^2) = \phi_M(s_n^2)$. Since $u(f^*) \geq \phi_M(s)$ and $\phi_M(s) \geq \phi_M(s_n^1)$ we obtain

$\phi_M(s_n^2) \geq \phi_M(s_n^1) \geq \phi_M(s_n^2)$ - implying that the value function ϕ_M is constant on the union of the cells $E_{i_m}^1 \cup E_{i_m}^2$.

For the inductive extension, take $\{E'_i\}$ to be the partition generated by $\mathcal{L}_1, \dots, \mathcal{L}_k$ and $m - 1$ of the unobservable hyperplanes in the set $\{\mathcal{L}_{k+1}, \dots, \mathcal{L}_n\}$. Denote the partition generated by the original k hyperplanes and the $m - 1$ unobservable ones as $\{E_i\}$. Add an additional unobservable hyperplane to the collection and let the corresponding partition be denoted as $\{\hat{E}_i\}$. Mimic the base step argument to show that a menu $M \in \Sigma_{\hat{E}_i}$ descends to a measurable function on the coarser partition, so that $M \in \Sigma_{E_i}$. Now apply the induction hypothesis to claim that, in fact, $M \in \Sigma_{E'_i}$. This shows that there are no observable restrictions to omitting hyperplanes which contain the state s_u and concludes the proof that there is a (canonical) consistent assignment rule on the boundary states - hence giving us a well-defined partition.

Returning to the proof of Theorem 2, the key step we need to prove is the full dimensionality of the set \mathcal{U}_ϕ in its vector span. Showing this full-dimensionality condition in the menu choice environment is the key technical innovation of the paper.

Descriptive Representations

For notational brevity I will identify menus M in Σ_{E_i} with elements of \mathcal{U}_ϕ , hence ignoring the distinction between a menu and its equivalence class under the map $M \mapsto \phi_M$. Let $\{E_i\}$ be the given Borel partition and let $\{(p_i, 1/K)\}$ be the associated lottery representation. Consider the set of pairs (p_i, p'_i) consisting of the lottery and its companion. Say that a pair of cells (E_i, E_j) in the Borel partition $\{E_i\}$ are *adjacent* if there is some lottery $(p_i^*, 1/K)$ that defines the partition such that the set $E_i \cup E_j$ is a cell in the Borel partition obtained by omitting this lottery. That is, $E_i \cup E_j$ is a cell in the Borel partition generated by the set of hyperplanes given by $\{(p_i, 1/K)\}_i \setminus (p_i^*, 1/K)$. I say a menu $M \in \mathcal{U}_\phi$ is a *descriptive* representation of the given Borel partition if it has the following properties,

1. $M = \{f_1, \dots, f_k\}$, where $f_i = \sum_{j=1}^k a_j p^i(j), p^i(j) \in \{p_j, p'_j\}$
2. $0 \leq a_j \leq 1$
3. $f_i = \arg \max_{f \in M} u_s(f), \forall s \in \text{int}(E_i)$
4. $|\{j : p^i(j) \neq p^{i'}(j)\}| = 1$ for all pairs $(f_i, f_{i'})$ where $(E_i, E_{i'})$ are adjacent.

I claim that every Borel partition has a menu which is a descriptive representation of the partition. Check this by induction on the number of hyperplanes that generate the Borel partition, i.e. the size of the set $\{(p_i, 1/K)\}$. Note that the result is obvious when this set is singleton. Inductively assume we can construct a descriptive representation M whenever the set has size $k - 1$ or less. Now let $\{(p_i, 1/K)\}_{i=1}^k$

denote a lottery representation of a Borel partition generated by k hyperplanes. Note that this is obtained as the smallest refinement of the partition generated by $\{(p_i, 1/K)\}_{i=1}^{k-1}$ and the partition generated by $\{(p_k, 1/K)\}$. Let M_{k-1} be a descriptive representation of the partition generated by $\{(p_i, 1/K)\}_{i=1}^{k-1}$ and define a new menu M_k as follows. Let $\{(a_1, \dots, a_{k-1}); (f_1, \dots, f_{k-1})\}$ denote M_{k-1} in shorthand. Let $\{E_k(1), E_k(2)\}$ denote the partition of the state space induced by the lottery representation $\{(p_k, 1/K)\}$, i.e. $E_k(1) := \{s \in \mathcal{S} : u_s(p_k) \geq u_s(p'_k)\}$, $E_k(2) := \{s \in \mathcal{S} : u_s(p_k) < u_s(p'_k)\}$. Also let $\{E'_i\}$ denote the Borel partition induced by the set $\{(p_i, 1/K)\}_{i=1}^{k-1}$. Put these partition cells into groups as follows. Define

$$\Sigma_i := \{E'_j : \text{int}(E'_j) \subseteq E_k(i)\}$$

Pick any constants b_1, b_2 such that $b_i > 0$ and $b_1 + b_2 = 1$ and for each cell E'_j such that $E'_j \in \Sigma_1 \cup \Sigma_2$ define an associated lottery \hat{f}_j as follows.

$$\hat{f}_j = \begin{cases} b_1 \cdot f_j + b_2 \cdot p_k, & \text{if } E'_j \in \Sigma_1 \\ b_1 \cdot f_j + b_2 \cdot p'_k, & \text{if } E'_j \in \Sigma_2 \end{cases}$$

For cells $E'_j \notin \Sigma_1 \cup \Sigma_2$ define \hat{f}_j as follows. For $s \in E'_j$ let $p_k(s) := \arg \max_{f \in \{p_k, p'_k\}} u_s(f)$. If the arg max is not singleton, let $p_k(s)$ denote the lottery with the higher $u(\cdot)$ -value. Note that for each $E'_j \notin \Sigma_1 \cup \Sigma_2$ we break up the cell into two new cells, $E'_j{}^1, E'_j{}^2$ where $\text{int}(E'_j{}^1) \subseteq E_k(1), \text{int}(E'_j{}^2) \subseteq E_k(2)$. Accordingly define adjusted lotteries \hat{f}_j as follows,

$$\hat{f}_j^i = \begin{cases} b_1 \cdot f_j + b_2 \cdot p_k, & \text{if } i = 1 \\ b_1 \cdot f_j + b_2 \cdot p'_k, & \text{if } i = 2. \end{cases}$$

Put

$$M_k := \bigcup_{j \in \Sigma_1 \cup \Sigma_2} \hat{f}_j \cup \bigcup_{j \notin \Sigma_1 \cup \Sigma_2} \{\hat{f}_j^1, \hat{f}_j^2\}$$

Notice that by the induction hypothesis (and by construction) the menu M_k is a descriptive representation of the Borel partition generated by the set $\{(p_i, 1/K)\}_{i=1}^k$. Therefore, all Borel partitions afford descriptive representations. I will now take a specific descriptive representation and show that one can generate a vector basis within \mathcal{U}_ϕ by taking suitable ‘perturbations’ of this representation.

The argument proceeds by induction on the number of hyperplanes that generate the partition $\{E_i\}$. The first step is to consider a special class of (coherent) Borel partitions - called *conical* partitions. These are partitions, $\{E_i\}$, with the following additional property. The intersection $\mathcal{L}_i \cap \mathcal{L}_j$ is identical for all pairs of hyperplanes $\mathcal{L}_i, \mathcal{L}_j$ that define the partition. It will be useful in what follows to consider the intersection $\mathcal{L}_i \cap \mathcal{L}_j$ in the full space \mathbf{R}^k , as opposed to just in the state

space \mathcal{S} . Note that this is wlog. In particular, for any 4-tuple (i, j, k, l) we have $\mathcal{L}_i \cap \mathcal{L}_j = \mathcal{L}_k \cap \mathcal{L}_l \Leftrightarrow \mathcal{L}_i \cap \mathcal{L}_j \cap \mathcal{S} = \mathcal{L}_k \cap \mathcal{L}_l \cap \mathcal{S}$. To see this check the \Leftarrow implication via contraposition. Say that there is some $s \in \mathbf{R}^k$ with $s \in \mathcal{L}_i \cap \mathcal{L}_j \setminus (\mathcal{L}_k \cap \mathcal{L}_l)$. Since $\mathcal{L}_i := \{s : (p_i - p'_i) \cdot s = 0\}$ (where each p_i is a lottery and p'_i is its companion), put $\vec{c} = [c, \dots, c]$ and note that $(p_i - p'_i) \cdot \vec{c} = 0$. Thus, find $c > 0$ large enough so that $(s + \vec{c})_i \geq 0, \forall i = 1, 2, \dots, k$ (with strictness for some i) and note that $(p_i - p'_i) \cdot s \geq (\leq) 0 \Leftrightarrow (p_i - p'_i) \cdot (s + \vec{c}) \geq (\leq) 0$. Put $\hat{s} := (s + \vec{c}) / \sum_{i=1}^k (s + \vec{c})_i$ to obtain that $\hat{s} \in \mathcal{L}_i \cap \mathcal{L}_j \setminus (\mathcal{L}_k \cap \mathcal{L}_l)$.

When $\mathcal{L}_i \cap \mathcal{L}_j = \mathcal{L}_k \cap \mathcal{L}_l$ for all (i, j, k, l) we call this a conical partition for the following reason. The intersection property implies that for any menu M which gives a descriptive representation $(M = \{(a_1, \dots, a_n); (p_{i_1}, \dots, p_{i_n})\})$ of the partition, the states that comprise, say, cell E_1 are described by

$$E_1 = \{s \in \mathcal{S} : (i) p_{i_1}(1)s \geq p_{i_1}(1)'s, \text{ and } (ii) p_{i_2}(1)s \geq p_{i_2}(1)'s\}$$

That is, E_1 is exactly the positive cone (in \mathcal{S}) formed by the two planes $p_{i_1}(1) - p_{i_1}(1)'$ and $p_{i_2}(1) - p_{i_2}(1)'$. For this reason, we call these conical Borel partitions. I now prove a claim for conical partitions that will be referred to as a ‘‘perturbation argument.’’ This claim will be invoked in the inductive extension. Label this as step 1 of the argument.

Step 1: Perturbation argument for conical partitions.

Let E_1, E_2, \dots, E_k be an enumeration of the cells of the partition, labeled so that the boundary of (E_i, E_{i+1}) is given by $\mathcal{L}_{i+1} := \{s : (p_{i+1} - p'_{i+1}) \cdot s = 0\}$. Assume the partition is generated by $\{\mathcal{L}_1, \mathcal{L}_2, \dots, \mathcal{L}_k\}$ and let $M = \{(a_1, a_2, \dots, a_k); (f_1, \dots, f_{2+2k})\}$ be a descriptive representation.²³ We label so that f_i is the lottery chosen by states in cell E_i . Now consider the following perturbed lotteries. Put

$$\hat{f}_1 := f_1 + \alpha_1 \cdot (p_1 - p'_1), \quad \hat{f}_2 := f_2 + \alpha_2 \cdot (p_3 - p'_3)$$

and, in addition, choose α_1, α_2 such that $u(\hat{f}_1) = u(\hat{f}_2)$. Let $\hat{M} := \{\hat{f}_1, \hat{f}_2, f_3, \dots, f_{2+2k}\}$. Consider the following hypothesis:

(H): There is a pair of states $(s_1, s_2) \in E_1 \times E_2$ with $\hat{f}_1 \succ_{s_1} \hat{f}_2, \hat{f}_2 \succ_{s_2} \hat{f}_1$.

Let us assume for the time being that I can choose perturbed lotteries \hat{f}_1, \hat{f}_2 such that hypothesis (H) holds. I now check that this implies that the menu \hat{M} is Σ_{E_i} -measurable. We then use this in step 2 to show that the span of the set \mathcal{U}_ϕ has full dimension.

²³The size of the menu, $2 + 2k$, turns out to be a consequence of the partition being conical, but this is an unimportant fact for what follows.

Consider the hyperplane $\mathcal{L} := \{s : \hat{f}_1 \cdot s = \hat{f}_2 \cdot s\}$. I claim that the partition obtained by adding the plane \mathcal{L} is also conical. To see this, label the planes that define the borders of cells E_1 and E_2 as p_n, p_1, p_2 . That is, p_n defines the border between cells E_n and E_1 , p_1 defines the border between cells E_1 and E_2 , etc. Note that we are making strong use of the conical structure here since we can associate cones formed by border planes with partition cells in a one-to-one manner. Let \mathcal{L}_i denote the plane corresponding to the border formed by the lottery p_i . I first claim that

$$\mathcal{L}_2 \cap \mathcal{L}_n \subseteq \mathcal{L} \cap \mathcal{L}_2$$

To see this note that $\mathcal{L}_2 \cap \mathcal{L}_n = \mathcal{L}_1 \cap \mathcal{L}_n = \mathcal{L}_1 \cap \mathcal{L}_2$. Therefore, if $s \in \mathcal{L}_2 \cap \mathcal{L}_n$, then $s \in \mathcal{L}$. The containment follows. Now notice that $\mathcal{L}_2 \cap \mathcal{L}_n$ and $\mathcal{L} \cap \mathcal{L}_2$ are both linear spaces of dimension one less than the dimension of the hyperplane. Since strict subspaces must have strictly smaller (linear) dimension, it follows that $\mathcal{L}_2 \cap \mathcal{L}_n = \mathcal{L} \cap \mathcal{L}_2$. Now observe that we have the following string of equalities,

$$\begin{aligned} \mathcal{L} \cap (\mathcal{L}_2 \cap \mathcal{L}_n) &= \mathcal{L}_2 \cap \mathcal{L}_n \\ \mathcal{L} \cap (\mathcal{L}_i \cap \mathcal{L}_j) &= \mathcal{L}_i \cap \mathcal{L}_j \quad (\text{since } \mathcal{L}_2 \cap \mathcal{L}_n = \mathcal{L}_i \cap \mathcal{L}_j) \end{aligned}$$

The latter implies that $\mathcal{L}_i \cap \mathcal{L}_j \subseteq \mathcal{L} \cap \mathcal{L}_i$. Comparing linear dimensions yields the equality $\mathcal{L} \cap \mathcal{L}_i = \mathcal{L}_i \cap \mathcal{L}_j$. Therefore, the augmented partition is also conical. Let $\{E_1, \dots, E_n\}$ denote the cells in the original partition and assume that E_i, E_j are such that $\mathcal{L} \cap \text{int}(E_i) \neq \emptyset$ and $\mathcal{L} \cap \text{int}(E_j) \neq \emptyset$. Let $\{(p_1 - p'_1), \dots, (p_n - p'_n)\}$ denote the set of normal vectors associated with the collection of hyperplanes, $\{\mathcal{L}_1, \dots, \mathcal{L}_n\}$, that form the partition. Let $\{(a_1, \dots, a_n); (f_1, \dots, f_n)\}$ be a descriptive representation of the partition, where $f_i = \sum_{j=1}^n a_j \cdot f_i(j)$ (and $f_i(j) \in \{p_j, p'_j\}$). Let $\text{supp}(f_i) := \{f_i(1), \dots, f_i(n)\}$. If \mathcal{L} intersects both $\text{int}(E_i)$ and $\text{int}(E_j)$, then I claim that $\text{supp}(f_i) \cap \text{supp}(f_j) = \emptyset$. Check this via contradiction. If there is some $f_* \in \text{supp}(f_i) \cap \text{supp}(f_j)$, find $g_* \notin \text{supp}(f_i) \cap \text{supp}(f_j)$ (which exists since $E_i \neq E_j$) and let $g_* \in \text{supp}(f_i), g'_* \in \text{supp}(f_j)$. Take $s_i \in \mathcal{L} \cap \text{int}(E_i)$ and $s_j \in \mathcal{L} \cap \text{int}(E_j)$ and note that

$$s_i \cdot (g_* - g'_*) > 0 \text{ and } s_j \cdot (g_* - g'_*) < 0$$

Find $\alpha \in (0, 1)$ so that $\alpha [s_i \cdot (g_* - g'_*)] + (1 - \alpha) [s_j \cdot (g_* - g'_*)] = 0$. Consider the state $s_\alpha := \alpha \cdot s_i + (1 - \alpha) s_j$ and note that $s_\alpha \cdot (g_* - g'_*) = 0$. Thus, letting \mathcal{L}_* denote the hyperplane in the collection $\{\mathcal{L}_1, \dots, \mathcal{L}_n\}$ with normal $(g_* - g'_*)$, we have $s_\alpha \in \mathcal{L}_*$. Since $s_\alpha \in \mathcal{L}$ we obtain, as the (augmented) partition is conical

$$s_\alpha \in \mathcal{L} \cap \mathcal{L}_* = \mathcal{L}_i \cap \mathcal{L}_j, \forall \mathcal{L}_i, \mathcal{L}_j \in \{\mathcal{L}_1, \dots, \mathcal{L}_n\}$$

Thus, $s_\alpha \cdot (p_i - p'_i) = 0, \forall i$. However, there is some $f_* \in \text{supp}(f_i) \cap \text{supp}(f_j)$, which implies (find p_k such that $f_* \in \{p_k, p'_k\}$) that either

1. $s_i \cdot (p_k - p'_k) > 0, s_j \cdot (p_k - p'_k) > 0$ or

$$2. s_i \cdot (p_k - p'_k) < 0, s_j \cdot (p_k - p'_k) < 0.$$

In either case, we obtain $s_\alpha \cdot (p_k - p'_k) \neq 0$ - contradiction.

It follows that if (E_i, E_j) are such that $\mathcal{L} \cap \text{int}(E_i) \neq \emptyset$ and $\mathcal{L} \cap \text{int}(E_j) \neq \emptyset$, then $\text{supp}(f_i) \cap \text{supp}(f_j) = \emptyset$. Now consider a state $s \in \mathcal{S}$ where, say, $\hat{f}_1 \in \arg \max_{f \in \hat{M}} u_s(f)$. Consider two cases: (i) $s \in \mathcal{L}$ and (ii) $s \notin \mathcal{L}$. In the first case, if $s \in \mathcal{L} \cap \mathcal{L}_k$ (for some k), then $s \in \mathcal{L}_i \cap \mathcal{L}_j, \forall (i, j)$ from which it follows that $\phi_{\hat{M}}(s) = \phi_M(s)$. Thus, consider $s \in \mathcal{L}, s \notin \mathcal{L}_i, \forall i$. If $s \notin E_1 \cup E_2$, then the preceding observation implies that $s \cdot (f_1(1) - f_1(1)') < 0$ (since (i) by hypothesis (H) there is a pair of interior states $(s_1, s_2) \in E_1 \times E_2$ with $\hat{f}_1 \succ_{s_1} \hat{f}_2$ and $\hat{f}_2 \succ_{s_2} \hat{f}_1$, and (ii) $s_i \cdot (f_1(1) - f_1(1)') > 0$ for all states in the interior of $E_1 \cup E_2$). But this implies that $s \cdot \hat{f}_1 < s \cdot f_1 \leq s \cdot f_i$ (where $f_i \in \arg \max_{f \in M} u_s(f)$) - contradiction. Thus, if $s \in \mathcal{L}$ and $\hat{f}_1 \in \arg \max_{f \in \hat{M}} u_s(\cdot)$, then $s \in E_1 \cup E_2$. Now consider the case where $s \notin \mathcal{L}$. Say that $\hat{f}_1 \succ_s \hat{f}_2$ with $\hat{f}_1 \in \arg \max_{f \in \hat{M}} u_s(\cdot)$. If $s \notin E_1 \cup E_2$ and, additionally, $s \cdot (f_1(1) - f_1(1)') = 0$, then we still have $\phi_{\hat{M}}(s) = \phi_M(s)$. Thus, we reduce to $\hat{f}_1 \succ_s \hat{f}_2$ and $f_1(1) \succ_s f_1(1)'$ (for $\hat{f}_1 \in \arg \max_{f \in \hat{M}} u_s(f)$). Label the hyperplanes in the original partition such that cells $E_1 \cup E_2$ (disregarding boundaries) are formed by the cone between \mathcal{L}_1 and \mathcal{L}_3 (hyperplane \mathcal{L}_i has normal $p_i - p'_i$). That is,

$$E_1 \cup E_2 = \{s \in \mathcal{S} : (p_1 - p'_1) \cdot s \geq 0, (p_3 - p'_3) \cdot s \geq 0\}$$

Also let $\mathcal{L} := \{s \in \mathcal{S} : (\hat{f}_1 - \hat{f}_2) \cdot s \geq 0\}$. Let

$$\mathcal{C}(\mathcal{L}_1, \mathcal{L}) := \{s \in \mathcal{S} : (p_1 - p'_1) \cdot s \geq 0, (\hat{f}_1 - \hat{f}_2) \cdot s \geq 0\}$$

I claim that $\mathcal{C}(\mathcal{L}_1, \mathcal{L}) \subseteq E_1 \cup E_2$. Check this via contradiction. Take an alleged $s_* \in \mathcal{C}(\mathcal{L}_1, \mathcal{L}) \setminus E_1 \cup E_2$ and find $s^* \in \text{int}(E_1 \cup E_2) \cap \text{int}(\mathcal{C}(\mathcal{L}_1, \mathcal{L}))$ (note that there is at least one such state by construction of \hat{f}_1, \hat{f}_2). Note that we have

1. $s_* \cdot (p_3 - p'_3) < 0$
2. $s^* \cdot (p_3 - p'_3) > 0$

Thus, find an $\alpha \in (0, 1)$ and $s_\alpha := \alpha \cdot s_* + (1 - \alpha) \cdot s^*$ such that $s_\alpha \cdot (p_3 - p'_3) = 0$, i.e. $s_\alpha \in \mathcal{L}_3 := \{s : (p_3 - p'_3) \cdot s = 0\}$. Note also that $s_\alpha \in \text{int}(\mathcal{C}(\mathcal{L}_1, \mathcal{L}))$. In particular, $s_\alpha \notin \mathcal{L}_1 \cap \mathcal{L}_3$. Consider the vector span, $V = \text{span}(\mathcal{L}_3 \cap \mathcal{L}_1, s_\alpha)$. Note that (i) $\dim(V) = k - 1$ and (ii) since $\mathcal{L}_3 \cap \mathcal{L}_1 = \mathcal{L} \cap \mathcal{L}_1$ (as the augmented partition is conical) we have $V \subseteq \mathcal{L}_3$.²⁴ Comparing linear dimensions we get $V = \mathcal{L}_3$. Since $V \subseteq \mathcal{C}(\mathcal{L}_1, \mathcal{L})$ this implies

$$E_1 \cup E_2 = \mathcal{C}(\mathcal{L}_3, \mathcal{L}_1) \subseteq \mathcal{C}(\mathcal{L}, \mathcal{L}_1)$$

²⁴Recall our earlier comment that allowed us to pass between the sets \mathcal{L} (which is a linear space) and $\mathcal{L} \cap \mathcal{S}$.

This contradicts the hypothesis (H) that there is some state $s \in E_1 \cup E_2$ for which $\hat{f}_2 \succ_s \hat{f}_1$. It follows that the only states $s \in \mathcal{S}$ for which $\hat{f}_1 \in \arg \max_{f \in \hat{M}} u_s(f)$ are states in the union of the cells E_1 and E_2 . This completes the proof that, for conical partitions, the perturbed menus \hat{M} are also measurable w.r.t. to the original Borel partition $\{E_i\}$.

Step 2: Full-dimensionality of non-conical partitions.

We proceed by induction on the number of hyperplanes that define the Borel partition. We first require an intermediate step where the partition is (conical) generated by two hyperplanes. We could simply invoke step 1 at this point and move on to non-conical partitions with three or more hyperplanes; however, we feel it is easier to appreciate some of the subtleties in the extension argument after going through the 2-hyperplane case in detail first. With this case in hand, I then show the full-dimensionality claim for non-conical partitions generated by 3 hyperplanes. The general inductive extension is easy to follow from this point on.

Sub-Step 2a: Full-dimensionality of two-hyperplane partitions.

To see this, let $\{\mathcal{L}_1, \mathcal{L}_2\}$ denote the two hyperplanes that generate the partition $\{E_i\}$, where we put $\mathcal{L}_1 := \{s \in \mathcal{S} : (p_1 - p'_1)s = 0\}$, $\mathcal{L}_2 := \{s \in \mathcal{S} : (p_2 - p'_2)s = 0\}$ and, by hypothesis, $p_i \sim_{u_s} p'_i$. Now add a hyperplane, \mathcal{L}^* , to this partition with the following two properties.

1. $\mathcal{L}^* = \{s : (p_* - p'_*)s = 0\}$, where $u(p_*) = u(p'_*)$.
2. $\mathcal{L}^* \cap \mathcal{L}_1 = \mathcal{L}^* \cap \mathcal{L}_2 = \mathcal{L}_1 \cap \mathcal{L}_2$.

That is, the partition of \mathcal{S} induced by these three hyperplanes is conical and, moreover, the state s_u (where the state-dependent vNM utility is equivalent to the normative ranking on lotteries) lies on the plane \mathcal{L}^* . It is straightforward to see that such a plane \mathcal{L}^* exists.²⁵ To distinguish this partition from the original two-hyperplane partition, let the augmented partition be denoted as $\{\overline{E}_i\}$ and the original partition as $\{E_i\}$. Let M be a descriptive representation of \overline{E}_i . Importantly, note that $M \in \Sigma_{E_i}$ (the original two-hyperplane partition). Let (E_1, E_2, E_3, E_4) be an enumeration of the cells in this original partition, and take $s_u \in E_1$. Take $V = \text{span}(\mathcal{L}_1 \cap \mathcal{L}_2, s_u)$ and note that V is a hyperplane. Denote this as \mathcal{L}_u . Consider adding this (unobservable) hyperplane to the original partition, so that we have a conical partition formed by $\{\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_u\}$. Coordinatize the cells for the partition formed by $\{\mathcal{L}_1, \mathcal{L}_2\}$ as follows,

1. $E_1 = \{s : (p_1 - p'_1) \cdot s \geq 0, (p_2 - p'_2) \cdot s \geq 0\}$
2. $E_2 = \{s : (p_1 - p'_1) \cdot s \geq 0, (p_2 - p'_2) \cdot s \leq 0\}$

²⁵To see this, let $W = \mathcal{L}_1 \cap \mathcal{L}_2$ and let s_u be the state with cardinal preference given by $u(\cdot)$. Consider $V = \text{span}(W, s_u)$ and note that V has linear dimension $k - 1$, hence is itself a hyperplane.

3. $E_3 = \{s : (p_1 - p'_1) \cdot s \leq 0, (p_2 - p'_2) \cdot s \leq 0\}$
4. $E_4 = \{s : (p_1 - p'_1) \cdot s \leq 0, (p_2 - p'_2) \cdot s \geq 0\}$.

Since $s_u \in \text{int}(E_1)$, the plane \mathcal{L}_u intersects E_1 and (by the argument in step 1) E_3 . Let $p_u - p'_u$ denote a normal to \mathcal{L}_u and coordinatize the cells of the partition formed by $\{\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_u\}$ as follows.

1. $E_1^1 = \{s : s \in E_1, (p_u - p'_u) \cdot s \geq 0\}$
2. $E_1^2 = \{s : s \in E_1, (p_u - p'_u) \cdot s \leq 0\}$
3. $E_3^1 = \{s : s \in E_4, (p_u - p'_u) \cdot s \geq 0\}$
4. $E_3^2 = \{s : s \in E_4, (p_u - p'_u) \cdot s \leq 0\}$

Let $M = \{(a_1, a_2, a_3); (f_1, \dots, f_6)\}$ be a descriptive representation of this partition and note that $\phi_M(s) = \phi_M(s'), \forall s, s' \in E_1$ and, similarly, $\phi_M(s) = \phi_M(s'), \forall s, s' \in E_3$ - so that the ϕ_M is measurable w.r.t. to the original partition $\{E_i\}$. I now construct a perturbed menu \hat{M} as in step 1 and verify that hypothesis (H) is satisfied for this menu. Once we show this, the argument from step 1 shows that $(1, 0, 0, 0)$ (in the $\{E_i\}$ coordinates) is in the span of \mathcal{U}_ϕ . Find states $s_1, \in E_1^1, s_2 \in E_1^2$ such that $p_u \succ_{s_1} p'_u$ and $p_u \succ_{s_2} p'_u$. For each s_i , find a ball $B_\epsilon(p_u), B_\epsilon(p'_u)$ such that (resp.) $p \succ_{s_1} p'_u, \forall p \in B_\epsilon(p_u)$ (and similarly for p'_u). Now find $\alpha_1, \alpha_2 > 0$ such that the following two properties are satisfied:

1. $p_u + \alpha_1 \cdot (p_1 - p'_1) \in B_\epsilon(p_u), p'_u + \alpha_2 \cdot (p_2 - p'_2) \in B_\epsilon(p'_u)$
2. $u(\hat{f}_1) = u(\hat{f}_2)$.

It follows that (H) is satisfied for this choice of α_i . Therefore, the span of \mathcal{U}_ϕ has dimension 4 - so that two-hyperplane partitions are full-dimensional.

Now apply this argument to the case where the partition consists of three planes $\{\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3\}$ where \mathcal{L}_3 is *transversal* to the two-hyperplane partition formed by \mathcal{L}_1 and \mathcal{L}_2 - that is, the plane \mathcal{L}_3 intersects the interiors of every cell in the two-hyperplane partition. The base step of our induction argument shows full-dimensionality of this class of Borel partitions.

Base Step: Full dimensionality of 3-hyperplane partitions (that are not conical). Coordinatize the cells as follows. Let $\{E_1, E_2, E_3, E_4\}$ denote the cells in the two-plane partition formed by \mathcal{L}_1 and \mathcal{L}_2 . Since the plane \mathcal{L}_3 is transverse to all of these cells, put $\mathcal{L}_3 := \{s \in \mathcal{S} : (p_3 - p'_3) \cdot s = 0\}$ and let $E_5 := E_1 \cap \{s : (p_3 - p'_3) \cdot s < 0\}, E_6 := E_2 \cap \{s : (p_3 - p'_3) \cdot s < 0\}$, and so on. In writing this coordinatization, I am implicitly assuming that \mathcal{L}_3 intersects the interior of all cells $\{E_1, E_2, E_3, E_4\}$ that

form the two-hyperplane partition created by \mathcal{L}_1 and \mathcal{L}_2 . To see this note that since the partition generated by $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$ is not conical, we know that $\mathcal{L}_1 \cap \mathcal{L}_2 \not\subseteq \mathcal{L}_3$. Now for a set V , let $T(V)$ denote the set of all half-spaces that contain V , where the half-spaces take the form $\{s \in \mathbf{R}^k : (p - p')s \geq (\leq) 0\}$ - the pair (p, p') consists of a lottery and its companion pair. Towards contradiction, say that \mathcal{L}_3 does not intersect the interior of E_1 and, abusing notation, let $E_1 = \overline{E_1}$ (the closure of E_1). Note that some boundary states need to be assigned to $\text{int}(E_1)$ under the assignment rule $\kappa(\cdot)$, so that E_1 is not necessarily closed. This distinction makes no difference for the forthcoming argument. Notice that we have the equality of sets

$$T(E_1) = T(\text{int}(E_1))$$

Let $H(\mathcal{L}, E_1) \in T(E_1)$ denote a half-space (of the given form) that contains E_1 . We then obtain

$$E_1 = \bigcap_{H \in T(E_1)} H(\mathcal{L}, E_1) = \bigcap_{H \in T(\text{int}(E_1))} H(\mathcal{L}, \text{int}(E_1))$$

If \mathcal{L}_3 does not intersect the interior of cell E_1 , then we know that either $\{s : (p_3 - p'_3) \cdot s \geq 0\} \supseteq \text{int}(E_1)$ or $\{s : (p_3 - p'_3) \cdot s \leq 0\} \supseteq \text{int}(E_1)$. Say the former holds. Then, the preceding sentence implies

$$\mathcal{L}_1 \cap \mathcal{L}_2 \subseteq \{s : (p_3 - p'_3) \cdot s \geq 0\}$$

However, $\mathcal{L}_1 \cap \mathcal{L}_2$ is a linear space - so that we must have $\mathcal{L}_1 \cap \mathcal{L}_2 \subseteq \mathcal{L}_3$. This contradicts the hypothesis that the partition formed by $\{\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3\}$ is not conical. Therefore, \mathcal{L}_3 has non-trivial intersection with $\text{int}(E_1)$, a symmetric argument applies to all other E_i . This shows that the chosen coordinatization is well-defined. I now claim that the rows of the following matrix are in the span of the set \mathcal{U}_ϕ (the set of utility vectors induced by menus in Σ_{E_i} - the 3-hyperplane partition),

$$A = \left[\begin{array}{cccc|cccc} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ \hline 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \end{array} \right]$$

Note that we are applying the full-dimensionality of two-hyperplane partitions in two separate places to claim that the rows of A are in the span of \mathcal{U}_ϕ . First, to claim that the top 4×8 block is in the span - here we are applying the full dimensionality of the partition formed by \mathcal{L}_1 and \mathcal{L}_2 . Second, to show that the bottom three rows are in the span of \mathcal{U}_ϕ we are applying the full-dimensionality claim to (resp.) the 2-plane partitions formed by $(\mathcal{L}_1, \mathcal{L}_3)$ and $(\mathcal{L}_2, \mathcal{L}_3)$ - since each of the rows correspond

to cells in one of these partitions. To show that $\text{span}(\mathcal{U}_\phi)$ has full dimension, I must produce one more row vector independent of the others in the matrix A . To do this, mimic the idea from sub-step 2a of adding an unobservable hyperplane to the partition. Let $\mathcal{L}_1, \mathcal{L}_2$ be conical and add a hyperplane \mathcal{L}_u with $s_u \in \mathcal{L}_u$ and such that $\{\mathcal{L}_u, \mathcal{L}_1, \mathcal{L}_2\}$ is conical. Now consider adding \mathcal{L}_3 to this partition. Since the plane \mathcal{L}_u is unobservable the partition formed by the planes $\{\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3, \mathcal{L}_u\}$ admits the same coordinatization as given above. In particular, we choose coordinates such that

1. $E_1 = \{p_1 \succ_s p'_1, p_2 \succ_s p'_2, p_3 \succ_s p'_3\}$
2. $E_2 = \{p_1 \succ_s p'_1, p_2 \prec_s p'_2, p_3 \succ_s p'_3\}$
3. $E_3 = \{p_1 \prec_s p'_1, p_2 \prec_s p'_2, p_3 \succ_s p'_3\}$
4. $E_4 = \{p_1 \prec_s p'_1, p_2 \succ_s p'_2, p_3 \succ_s p'_3\}$
5. $E_5 = \{p_1 \succ_s p'_1, p_2 \succ_s p'_2, p_3 \prec_s p'_3\}$
6. $E_6 = \{p_1 \succ_s p'_1, p_2 \prec_s p'_2, p_3 \prec_s p'_3\}$
7. $E_7 = \{p_1 \prec_s p'_1, p_2 \prec_s p'_2, p_3 \prec_s p'_3\}$
8. $E_8 = \{p_1 \prec_s p'_1, p_2 \succ_s p'_2, p_3 \prec_s p'_3\}$.

We label the vectors (p_i, p'_i) such that $p_i \succ_{s_u} p'_i$. Take $M := \{(a_1, a_9); (f_1, f_2, \dots, f_{12})\}$ to be descriptive representation of the partition formed by the planes $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$ and the unobservable hyperplane \mathcal{L}_u . Since \mathcal{L}_u intersects cells E_1 and E_5 we choose the labeling in the descriptive representation such that f_1, f_2 are chosen by states in cell E_1 and f_7 and f_8 are chosen by states in cell E_5 . Now consider a perturbed menu \hat{M} as follows. As in the text put $\hat{f}_7 = f_7 + \alpha_1 \cdot (p_1 - p'_1)$, $\hat{f}_8 = f_8 + \alpha_2 \cdot (p_2 - p'_2)$. Consider the hyperplane $\mathcal{L}^* = \{s : (\hat{f}_7 - \hat{f}_8) \cdot s = 0\}$. The argument in step 1 implies that the partition formed by $\{\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_u, \mathcal{L}^*\}$ is conical. Hence, exactly as in step 1, we obtain the inclusion

$$\mathcal{C}(\mathcal{L}^*, \mathcal{L}_1) \subseteq \mathcal{C}(\mathcal{L}_1, \mathcal{L}_2)$$

This implies that the only states $s \in \mathcal{S}$ for which either (i) \hat{f}_7 or $\hat{f}_8 \in \arg \max_{\hat{M}} u_s(\cdot)$ are states in the union of the cells $E_1 \cup E_5$ (ignoring boundaries). Since $u(p_1 - p'_1), u(p_2 - p'_2) > 0$ we can choose α_1, α_2 such that $u(f_1) = u(f_2) = u(\hat{f}_7) = u(\hat{f}_8)$ - that is, the function $\phi_{\hat{M}}(\cdot)$ is *constant* on the union $E_1 \cup E_2$. Since $\phi_M = \phi_{\hat{M}}$ outside of $E_1 \cup E_2$, it follows that $\phi_{\hat{M}}$ is measurable w.r.t. the partition formed by $\{\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3\}$. Moreover, note that $\phi_{\hat{M}} - \phi_M$ is proportional to $[0, 0, 0, 0 | 1, 0, 0, 0]$. This implies that $[1, 0, 0, 0 | 0, 0, 0, 0]$ is in the span of \mathcal{U}_ϕ . Thus, the span of \mathcal{U}_ϕ

contains the rows of the following matrix,

$$\left[\begin{array}{cccc|cccc} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ \hline 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right]$$

It is straightforward to see that the above matrix has full rank. Take this to be the base step of the induction argument. Now we extend to partitions generated by an arbitrary (but finite) number of hyperplanes. Consider a partition $\{E_i\}$ formed by planes $\{\mathcal{L}_1, \dots, \mathcal{L}_n\}$. Label so that $\{\mathcal{L}_1, \dots, \mathcal{L}_k\}$ forms a maximal conical partition and such that $s_u \in \mathcal{C}(\mathcal{L}_1, \mathcal{L}_2)$. Let E_1, E_2, \dots, E_m be an enumeration of cells of the partition, where E_i, E_{i+1} are adjacent cells in a descriptive representation.²⁶ Note that, by the induction hypothesis, the rows of the following matrix are all in the span of \mathcal{U}_ϕ .

$$\left[\begin{array}{ccccccc} 1 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & \cdots & 0 & 1 & 1 \end{array} \right]_{n-1 \times n}$$

Note that cells E_1, E_2 are separated by a single plane, say $\mathcal{L}_k := \{s : (p_k - p'_k) \cdot s = 0\}$. Thus, once we remove this plane the union $E_1 \cup E_2$ is a cell in the coarsened partition (call it $\{E'_i\}$) obtained by removing \mathcal{L}_k . By the induction hypothesis, the vector $[1, 1, 0, \dots, 0]$ is in $\text{span}(\mathcal{U}_\phi \cap \Sigma_{E'_i})$ where we identify $\Sigma_{E'_i}$ with its (embedded) image in Σ_{E_i} . Similarly argue to show that the remaining rows of A are in the span of \mathcal{U}_ϕ . We now add an unobservable hyperplane to the partition. Arguing as in the base step, add \mathcal{L}_u to the partition, where $\mathcal{L}_u = \text{span}(s_u, \mathcal{L}_1 \cap \mathcal{L}_2)$ (here we take $s_u \in \mathcal{C}(\mathcal{L}_1, \mathcal{L}_2)$). Let E_1, \dots, E_k be a list of the cells that are contained in the cone $\mathcal{C}(\mathcal{L}_1, \mathcal{L}_2)$ and let f_1, \dots, f_k be the lotteries (resp.) selected by states in E_1, \dots, E_k (from the descriptive representation M), and let $f_i^u, f_i^{u'}$ denote the lotteries with small weights attached (resp.) on the pair (p_u, p'_u) , where $p_u - p'_u$ is a normal vector to the unobservable hyperplane. Consider the following perturbed menu, \hat{M} . Put $\hat{f}_1^1 = f_1^u + \alpha_1^1(p_1 - p'_1), \hat{f}_1^2 = f_1^{u'} + \alpha_1^2(p_2 - p'_2), \dots, \hat{f}_k^1 = f_k^u + \alpha_k^1(p_1 - p'_1), \hat{f}_k^2 = f_k^{u'} + \alpha_k^2(p_2 - p'_2)$. Leave the lotteries $f_i, \forall k + 1 \leq i \leq m$ chosen by states in cells

²⁶If f_i is chosen by states in cell E_i and f_{i+1} is chosen by states in cell E_{i+1} , then the adjacency condition is that there is exactly one vector p_k such that $p_k \in \text{supp}(f_i)$ and $p'_k \in \text{supp}(f_{i+1})$. For all other vectors pairs (p_j, p'_j) we have either $p_j \in \text{supp}(f_i) \cap \text{supp}(f_{i+1})$ or $p'_j \in \text{supp}(f_i) \cap \text{supp}(f_{i+1})$.

$E_i, \forall i \geq k+1$ unchanged and put $\hat{M} := \{(\hat{f}_i^1, \hat{f}_i^2)_{i=1}^k, f_{k+1}, \dots, f_m\}$. I claim that (i) \hat{M} is Σ_{E_i} -measurable and (ii) we can choose the pairs (α_i^1, α_i^2) such that $\phi_{\hat{M}}$ is constant on the cone $\mathcal{C}(\mathcal{L}_1, \mathcal{L}_2)$. The second claim mimics the prior arguments verbatim - First, adjust the weights on the descriptive representations so that the differences $\phi_M(s) - \phi_M(s')$ are very small for all $s, s' \in \mathcal{C}(\mathcal{L}_1, \mathcal{L}_2)$. Next, choose α_i^1, α_i^2 such that $u(\hat{f}_i^1) = u(\hat{f}_i^2) = u(\hat{f}_j^1) = u(\hat{f}_j^2), \forall (i, j)$. For the Σ_{E_i} -measurability claim simply note that if \hat{f}_i^1 or \hat{f}_i^2 is selected (say \hat{f}_i^1) from \hat{M} for any alleged state outside the cone $\mathcal{C}(\mathcal{L}_1, \mathcal{L}_2)$, then we may consider the plane $\mathcal{L} := \{s : (\hat{f}_i^1) - \hat{f}_i^2 \cdot s \geq 0\}$ and apply the argument (via contradiction) from step 1. Namely, show that (by choice of the constants α_i) we have the inclusion of cones $\mathcal{C}(\mathcal{L}, \mathcal{L}_1) \subseteq \mathcal{C}(\mathcal{L}_1, \mathcal{L}_2)$. Having checked Σ_{E_i} measurability, consider the difference of the two functions, $\phi_{\hat{M}} - \phi_M$. Wlog say that we have the ordering $\phi_M(s_1) > \phi_M(s_2) > \dots > \phi_M(s_k), s_i \in \text{int}(E_i)$. We thus obtain $\phi_{\hat{M}} - \phi_M = [0, r_2, r_3, \dots, r_k, 0, 0, \dots, 0], r_i < r_{i+1}$. This then implies that

$$[0, 0, r_3 - r_2, r_4, \dots, r_k, 0, 0, \dots, 0] \in \text{span}(\mathcal{U}_\phi)$$

which in turn implies that $[0, 0, 0, r_4 - (r_3 - r_2), r_5, \dots, r_k, 0, 0, \dots, 0] \in \text{span}(\mathcal{U}_\phi)$. Inductively proceeding we obtain that $[0, 0, \dots, r_k - \sum_{i=1}^{k-2} (-1)^{i-1} \cdot r_{k-i}, 0, 0, \dots, 0] \in \mathcal{U}_\phi$. It follows that $\text{span}(\mathcal{U}_\phi)$ is full-dimensional.

This completes the proof of Theorem 2. The forthcoming arguments concern the extension of this result to the domain of menus that are $\sigma(\mathcal{A})$ -measurable.²⁷ This part makes strong use of the duality between the cone in \mathbf{R}^k generated by the set of differences, $\{\phi_M - \phi_{M'} : \phi_M \succeq \phi_{M'}\}$ and the ‘dual’ cone (also in \mathbf{R}^k since Euclidean space is self-dual) of probability measures on the σ -algebra $\{E_i\}$. Since we will need to make frequent reference to these objects, we label them as follows. Put

$$\mathcal{C}_{E_i} := \bigcup_{r>0} r\mathcal{H}, \quad \mathcal{H} := \{h \in \mathbf{R}^k : h = \phi_M - \phi_{M'}, M \succeq M'\}$$

Let $\mathcal{L}(\mathbf{R}^k)$ be the (vector) space of linear functionals on \mathbf{R}^k and denote the canonical ‘duality map’ via $t : \mathbf{R}^k \rightarrow \mathcal{L}(\mathbf{R}^k)$. Since the set \mathcal{C}_{E_i} has non-empty interior, axiom A4 (and the self-duality of \mathbf{R}^k) implies that $\mathcal{C}_{E_i}^t$ can be uniquely identified with a (closed and convex) set of probability measures (via the pullback map) on the state space \mathcal{S} .

²⁷This portion of the argument (from this sentence until the beginning of the construction of the inverse limit below. is directed towards the extension of Theorem 2, i.e. where the domain of the menu preference includes the σ -closure of the set \mathcal{A} . We have not yet succeeded in providing this extension, so that the remarks on directed sets, etc. are not relevant for the proof of Theorem 2 as currently stated. Thus, the reader may skip this portion of the argument without any loss in continuity.

the upwards directed system in X to be as follows. Let Λ be the collection of all (coherent) Borel partitions of \mathcal{S} and take the relation \succeq^* to be $\{E_i\} \succeq^* \{F_i\}$ iff $\{E_i\}$ is a refinement of $\{F_i\}$ - this is clearly a well-defined partial order. Consider the collection $\{\Sigma_{E_i}\}$, the system of E_i -measurable menus and let $\mathcal{C} := \{\mathcal{C}_{E_i}\}$ - where \mathcal{C}_{E_i} is the positive cone of $\{E_i\}$ measurable functions. Let $\mathcal{C}_{E_i}^t$ denote the set of probability measures on the Borel partition $\{E_i\}$, obtained from the finite subjective Bewley theorem. Similarly define $\mathcal{C}(B_{\text{isc}}^*(\mathcal{S}))$. Finally, let $\text{ca}(\mathcal{S})$ denote the set of (σ -finite) measures on the state space \mathcal{S} . Now notice that the system $\{\mathcal{C}_{E_i}\}$ is upwards directed and the system $\{\mathcal{C}_{E_i}^t\}$ is downwards directed. The former claim is obvious. The latter claim follows from an application of the finite Bewley theorem: Notice that if $\{F_i\} \succeq^* \{E_i\}$, then the set $\Sigma_{E_i} \subseteq \Sigma_{F_i}$ and is closed. Thus, $\mathcal{C}_{E_i} \subseteq \mathcal{C}_{F_i}$ (as a closed subset). Let $\iota(\mathcal{C}_{E_i})$ denote the image of \mathcal{C}_{E_i} in \mathcal{C}_{F_i} . For each $\pi \in \mathcal{C}_{F_i}^t$ we put

$$\text{proj}(\pi)(D) := \sum_{F_i: F_i \cap D \neq \emptyset} \pi(F_i)$$

where $D \in \{E_i\}$. Note that $\text{proj}(\pi)$ is a well-defined probability measure on $\{E_i\}$. Let x, y denote generic elements of Σ_{E_i} and note that we have the following sequence of implications,

$$x \succeq y \Leftrightarrow \iota(x) \succeq \iota(y)$$

$$\Leftrightarrow \int_{\{F_i\}} \iota(x) d\pi \geq \int_{\{F_i\}} \iota(y) d\pi, \quad \forall \pi \in \mathcal{C}_{F_i}^t$$

$$\Leftrightarrow \int_{\{E_i\}} x d\text{proj}(\pi) \geq \int_{\{E_i\}} y d\text{proj}(\pi), \quad \forall \pi \in \mathcal{C}_{F_i}^t.$$

Since this holds for all pairs $x, y \in \mathcal{C}_{E_i}$, by the uniqueness of the Bewley representation we obtain that the projection map, $\text{proj} : \mathcal{C}_{F_i} \rightarrow \mathcal{C}_{E_i}$, is a surjection. Uniqueness of the Bewley representation for Σ_{E_i} -measurable menus is standard, but in the interest of completeness we provide the formal argument here. Fix \mathbf{R}^n and let $C \subseteq \mathbf{R}^n$ be a closed and convex cone. Assume we have two sets of closed, convex functionals in $\mathcal{L}(\mathbf{R}^n, \mathbf{R})$ (call them $\mathcal{L}_1, \mathcal{L}_2$) such that all elements of $\mathcal{L}_1, \mathcal{L}_2$ are represented by dot products with probability measures and

$$C = \bigcap_{\ell \in \mathcal{L}_1} \{x : \ell(x) \geq 0\} = \bigcap_{\ell \in \mathcal{L}_2} \{x : \ell(x) \geq 0\}$$

Towards contradiction, say that $\mathcal{L}_1 \neq \mathcal{L}_2$. Take $\ell_1 \in \mathcal{L}_1 \setminus \mathcal{L}_2$. Note that $\{\ell_1\}$ and \mathcal{L}_2 are closed, convex, and disjoint sets in $\mathcal{L}(\mathbf{R}^n, \mathbf{R})$ (where $\{\ell_1\}$ is obviously bounded). By the Separating Hyperplane Theorem, there is a functional $\hat{x} \in \mathcal{L}(\mathcal{L}(\mathbf{R}^n, \mathbf{R}), \mathbf{R})$ such that $\hat{x}(\ell_1) < c < \hat{x}(\ell_2), \forall \ell_2 \in \mathcal{L}_2$. Since \mathbf{R}^n is reflexive, the functional \hat{x} is given

by point evaluation for some $x \in \mathbf{R}^n$. That is, $\hat{x}(\ell) = \ell(x), \forall \ell \in \mathcal{L}(\mathbf{R}^n, \mathbf{R})$ (for some $x \in \mathbf{R}^n$). Since the functionals in $\mathcal{L}_1, \mathcal{L}_2$ are all represented by dot products with probability measures, consider $x' := x - c \cdot \vec{1}$ and note that $\ell(x') = \ell(x) - c, \forall \ell \in \mathcal{L}_1, \mathcal{L}_2$. It follows that

$$\ell_1(x') < 0 < \ell_2(x'), \forall \ell_2 \in \mathcal{L}_2$$

Therefore, since $C = \cap_{\ell \in \mathcal{L}_2} \{x : \ell(x) \geq 0\}$, we obtain $x' \in C$. OTOH, we also have $C = \cap_{\ell \in \mathcal{L}_1} \{x : \ell(x) \geq 0\}$. Thus, if $x' \in C$ then $x' \in \{x : \ell_1(x) \geq 0\}$ - contradiction. It follows that $\mathcal{L}_1 = \mathcal{L}_2$.

Now we apply this to the sets of functionals at hand. Take $C' := \mathcal{C}_{F_i}, C := \mathcal{C}_{E_i}$, where $\{F_i\} \succeq^* \{E_i\}$ and, by abuse of notation, let $\mathcal{C}_{E_i}^t, \mathcal{C}_{F_i}^t$ denote the (closed and convex) collections of equivalence classes of the dual cone of linear functionals - so that for each class of functionals induced by the same probability measure we select the one such that $\ell_\pi(\vec{1}) = 1$. Note that, by the sufficiency argument for the representation (applied to Σ_{F_i}), we have

$$\mathcal{C}_{F_i} = \cap_{\ell \in \mathcal{C}_{F_i}^t} \{x : \ell(x) \geq 0\}$$

Now pull back both sides of this equality using the ι map. That is, the preceding equality yields: $\iota^{-1}(\mathcal{C}_{F_i}) = \iota^{-1}(\cap_{\ell \in \mathcal{C}_{F_i}^t} \{x : \ell(x) \geq 0\})$. Note that $\iota^{-1}(\mathcal{C}_{F_i}) = \mathcal{C}_{E_i}$ and

$$(*) \quad \iota^{-1}(\cap_{\ell \in \mathcal{C}_{F_i}^t} \{x : \ell(x) \geq 0\}) = \cap_{\text{proj}(\ell) : \ell \in \mathcal{C}_{F_i}^t} \{x : \text{proj}(\ell)(x) \geq 0\}$$

where the notation $\text{proj}(\ell)$ denotes the (normalized) linear functional represented by dot product with the probability measure $\text{proj}(\pi)$, where $\pi \in \mathcal{C}_{F_i}^t$ is the probability measure that represents ℓ . Thus, put $\text{proj}(\mathcal{C}_{F_i}^t)$ to be the set of projected functionals and put $\mathcal{L}_1 = \text{proj}(\mathcal{C}_{F_i}^t), \mathcal{L}_2 = \mathcal{C}_{E_i}^t$. By (resp.) $(*)$ and the Bewley representation applied to Σ_{E_i} we have

$$\mathcal{C}_{E_i} = \bigcap_{\ell \in \mathcal{L}_1} \{x : \ell(x) \geq 0\} = \bigcap_{\ell \in \mathcal{L}_2} \{x : \ell(x) \geq 0\}$$

Applying the argument in the previous paragraph we obtain $\mathcal{L}_1 = \mathcal{L}_2$ - proving uniqueness.

Returning to the main argument, it follows that the upwards directed system, $\mathcal{C} := \{\mathcal{C}_{E_i}\}$ corresponds, via the Bewley representation, to the downwards directed system, $\mathcal{C}^t := \{\mathcal{C}_{E_i}^t\}$. Now we show that (i) the inverse limit of the downwards system exists, call it $\varprojlim \mathcal{C}^t \subseteq \text{ca}(\mathcal{S})$, in the space $\text{ca}(\mathcal{S})$ and (ii) the topological dual of this system is exactly the cone, $\mathcal{C}(B_{\text{isc}}^*(\mathcal{S}))$. This will yield the Bewley representation for the full space of menus.

I first check that the inverse limit exists and is closed. Let $X := \prod_{\{E_i\}} \Delta(\{E_i\})$ and endow this space with the product topology (recall: the topology on $\text{ca}(\mathcal{S})$ is the topology of weak convergence). Consider the set $\overline{X} \subseteq X$ defined by $\overline{X} := \prod_{\{E_i\}} \mathcal{C}_{E_i}^t$. Let $\varprojlim \mathcal{C}^t$ be defined as follows. Since there are uncountably many (coherent) Borel partitions $\{E_i\}$ and this collection is endowed with a natural partial order \succ^* , let us index the collection with the set Λ . Put

$$\varprojlim \mathcal{C}^t := \{\pi = (\pi_\alpha, \pi_\beta, \dots) : \gamma \succ^* \alpha, \beta \Rightarrow \text{proj}_{\gamma, \alpha}(\pi_\gamma) = \pi_\alpha, \text{proj}_{\gamma, \beta}(\pi_\gamma) = \pi_\beta\}$$

Notice that this is clearly a closed subset of X . Moreover, notice that it can naturally be identified with a subset of the set of measures on \mathcal{A} . To see this, let \mathcal{B} denote the σ -algebra on \mathcal{S} generated by all (finite) Borel partitions of \mathcal{S} and recall that $\mathcal{A} = \bigcup_{\alpha \in \Lambda} \sigma(\{E_i\})$ is the algebra generated by all finite Borel partitions. Notice that every element of $\varprojlim \mathcal{C}^t$ determines (uniquely) a (finitely-additive) measure on the algebra \mathcal{A} . Similarly, any measure on \mathcal{A} determines an element of the set $\varprojlim \mathcal{C}^t$. In other words, as sets we can identify $\varprojlim \mathcal{C}^t$ with the collection of measures on \mathcal{A} . Since the product topology on the set $\varprojlim \mathcal{C}^t$ coincides with the topology of weak convergence²⁸ when we identify this set with measures on \mathcal{A} , the identification is a topological isomorphism.

²⁸Here we take the class of test functions to be the set of \mathcal{A} -measurable functions.

References

- [1] Bewley, T. (1986): “Knightian Decision Theory: Part I,” Cowles Foundation Discussion Papers.
- [2] Dekel, E., B. Lipman, and A. Rustichini (2001): “Representing Preferences with a Unique Subjective State Space,” *Econometrica*, vol. 69, 891 - 934.
- [3] Dubra, J. F. Maccheroni, and E. Ok (2004): “Expected Utility Theory without the Completeness Axiom,” *Journal of Economic Theory*, vol. 115, 118-133.
- [4] Gul, F. and W. Pesendorfer (2001): “Temptation and Self-Control,” *Econometrica*, vol. 69, 1403 - 1435.
- [5] Gul, F. and W. Pesendorfer (2005): “The Revealed Preference Theory of Changing Tastes,” *Review of Economic Studies*, vol. 72, 429 - 448.
- [6] Ok, E., P. Ortoleva, and G. Riella (2009): “Incomplete Preferences under Uncertainty: Indecisiveness in Beliefs vs. Tastes,” Mimeo, NYU and Caltech.
- [7] Stovall, J. (2009): “Multiple Temptations,” *Econometrica*, forthcoming.