NON-REDUNDANT BOREL SPACE OF LEXICOGRAPHIC EXPECTED UTILITY PREFERENCES*

Byung Soo Lee†

2012-02-19

Abstract

Lexicographic probability systems (LPS's) are representations of lexicographic expected utility (LEU) preferences, which were axiomitized by Blume, Brandenburger, and Dekel [1991]. LPS's that satisfy the *mutual singularity* condition can be viewed as providing beliefs conditional on events with zero prior probability à la Rényi [1955]. However, this interpretation loses much of its appeal when the space of uncertainty contains redundancies. The problem arises acutely in the construction of higher-order LEU preferences because the space of first-order lexicographic beliefs will contain uncountably many redundant representations of the same preference relation. It follows that the mutual singularity condition lacks bite when it is imposed on higher-order beliefs. In this paper, we resolve this issue by showing that there is a standard Borel space of LEU preferences without such redundancies.

*This is a *very* preliminary version of the paper.*

†E-mail: `byungsoolee@rotman.utoronto.ca` / Address: Rotman School of Management, University of Toronto, Toronto, ON M5S 3E6.

1 INTRODUCTION

Lexicographic probability systems (LPS's) can most simply be described as finite sequences of probability measures. They are representations of lexicographic expected utility (LEU) preferences, which were axiomitized by Blume, Brandenburger, and Dekel [1991]. LPS's that satisfy what is called the *mutual singularity condition* can be viewed as providing beliefs conditional on events with zero prior probability à la Rényi [1955]. Consider the following example, which is somewhat contrived but nevertheless informative:

$$
\begin{aligned}
X &= \{1,2,3,4\} & \qquad \mu_1(1) &= \mu_1(3) = 1/2 \\
\sigma &= (\mu_1,\mu_2) & \qquad \mu_2(2) &= \mu_1(4) = 1/2
\end{aligned}
$$

The LPS σ satisfies mutual singularity because the two component measures μ_1 and μ_2 give probability 1 to the disjoint events $Odd = \{1,3\}$ and $Even = \{2,4\}$, respectively. The primary belief μ_1 assigns probability 0 to the event *Even*. The secondary belief μ_2 can be viewed as the belief conditional on the event *Even*.

However, this natural interpretation loses much of its appeal when the space of uncertainty contains redundancies. In the following example, $1'$ is a redundant representation of 1.

$$
\begin{aligned}
X &= \{1,1',2,3\} & \qquad \mu_1(\{1\}) &= \mu_1(\{3\}) = 1/2 \\
\sigma &= (\mu_1,\mu_2) & \qquad \mu_2(\{2\}) &= \mu_2(\{1'\}) = 1/2
\end{aligned}
$$

The LPS σ satisfies mutual singularity as before, but the interpretation of μ_2 as the belief conditional on $\{1',2\}$ loses saliency when 1 is given positive probability by the primary belief μ_1.

The issue arises acutely in the construction of higher-order LEU preferences. The space $\mathcal{L}(X)$ of first-order lexicographic beliefs over X contains numerous redundant representations. For each LPS $\sigma \in \mathcal{L}(X)$, there exist an uncountable number of LPS's in $\mathcal{L}(X)$ that represent the exact same LEU preference relation. This holds true even when X itself contains no redunancies. It is therefore difficult to meaningfully impose the mutual singularity condition on second-order beliefs, which belong to the space $\mathcal{L}(X \times \mathcal{L}(X))$.

The axiom of choice guarantees the existence of a set $U \subseteq \mathcal{L}(X)$ that contains a representation of each and every LEU preference without redundancies. However, in order to define useful second-order preferences, the set U needs to have nice properties. For one thing, U would need to be a measurable set.

The mutual singularity condition plays an important role in the framework used by Brandenburger, Friedenberg, and Keisler [2008] to conduct epistemic analyses of iterated elimination of weakly dominated strategies. Each type in their framework maps to an LPS over the strategies and types of the other players. If two types represent the same hierarchy of beliefs, then it is not clear what modeling assumptions are captured by the mutual singularity condition. Given the aforementioned issues, a lexicographic type structure that contains all hierarchies of preferences cannot be constructed without redundancies unless we can pick out a well-behaved subset of $\mathcal{L}(X)$ for each arbitrary X.

In this paper, we show that there is indeed such a well-behaved non-redundant space of all LEU preferences represented in $\mathcal{L}(X)$ when X is a Polish space. In particular, it is a standard Borel space generated by a Polish topology, which means that the space of higher-order LEU preferences is always Polish as well. A similar property, which holds for higher-order expected utility preferences, is crucial to showing the existence of a single probability measure that extends all finite-order beliefs in a consistent infinite hierachy.

The paper is organized as follows. Section 2 gives a review of the basic terminology that is necessary to formally state our main problem and its solution, which borrows extensively from the work on Borel cross-section problems in the field of descriptive set theory. Section 3 breaks the proof of the main theorem into several substantially distinguishable sub-problems and provides a concise framework for discussing maximally parsimonious representations of LEU preference relations.

2 PRELIMINARIES

2.1 LEXICOGRAPHIC PROBABILITY SYSTEMS

Definition 1. *A* ***topological space*** *is a pair $(X, \mathcal{T})$, where X is a set and $\mathcal{T}$ is a topology on X.*

For the sake of brevity, we will often refer to topological spaces without specifying the topology. For example, when we refer to a topological space X, it should be implicitly understood that X is associated with some topology, which we will denote by $\mathbf{T}(X)$.

Definition 2. *A* ***Borel space*** *is a pair $(X, \mathcal{B})$, where X is a set and $\mathcal{B}$ is a Borel σ-algebra generated by some topology on X.*

For the sake of brevity, we will often refer to Borel spaces without specifying the Borel σ-algebra. For example, when we refer to a Borel space X, it should be implicitly understood that X is associated with some Borel σ-algebra, which we will denote by $\mathbf{B}(X)$. With some abuse of notation, we will also denote the Borel σ-algebra generated by the topology $\mathcal{T}$ as $\mathbf{B}(\mathcal{T})$.

Naturally, each topological space $(X, \mathcal{T})$ canonically generates the associated Borel space $(X, \mathbf{B}(\mathcal{T}))$. As such, if we declare at some point that X is a topological space then we may subsequently let X denote the associated Borel space as well when there is no risk of confusion.

Definition 3. *A topological space $(X, \mathcal{T})$ is called* ***Polish*** *if its topology $\mathcal{T}$ is separable and admits a complete metric.*

Definition 4. *A* ***standard Borel space*** *is Borel space $(X, \mathcal{B})$ where $\mathcal{B}$ is generated by a Polish topology on X.*

Definition 5. *Let X be a Polish space. A* ***lexicographic probability system (LPS)*** *on X is a finite sequence of probability measures on X. The space of all LPSs on X is denoted by $\mathcal{L}(X)$. The space of all probability measures on X is denoted by $\mathcal{P}(X)$.*

Definition 6. *The LPS $\sigma \in \mathcal{L}(X)$ has* ***length*** *equal to $n \in \mathbb{N}$—i.e., is length-n—if $\sigma = (\mu_1, \ldots, \mu_n)$ for some $\mu_1, \ldots, \mu_n \in \mathcal{P}(X)$. If σ is length-n then we write $\#\sigma = n$. The set of all length-n LPSs on X is denoted by $\mathcal{L}_n(X)$. It is readily seen that the following equations also define the sets $\mathcal{L}(X)$ and $\mathcal{L}_n(X)$ as well as the space $\mathcal{L}_{\leq n}(X)$ of all LPSs that have length less than or equal to n.*

$$\mathcal{L}_n(X) \equiv \prod_{k=1}^{n} \mathcal{P}(X); \qquad \mathcal{L}(X) \equiv \bigcup_{n \in \mathbb{N}} \mathcal{L}_n(X); \qquad \mathcal{L}_{\leq n}(X) \equiv \bigcup_{k=1}^{n} \mathcal{L}_k(X).$$

2.2 LEXICOGRAPHIC EXPECTED UTILITY PREFERENCES

Definition 7. *Let X be a Borel space. An* ***act*** *defined over X is a Borel map from X to $[0, 1]$. The set of all acts defined over X is denoted by $\mathcal{A}(X)$.*

Definition 8. *For each LPS $\sigma \in \mathcal{L}(X)$, $\succsim_\sigma$ is the preference relation over $\mathcal{A}(X)$ that is defined by the following sentence, where $\geq^{\mathcal{L}}$ denotes the usual lexicographic comparison over vectors of real numbers.*[1]

$$\begin{aligned} &\forall \sigma = (\mu_1, \ldots, \mu_n) \in \mathcal{L}(X), \forall f, g \in \mathcal{A}(X), \\ &\qquad f \succsim_\sigma g \iff \left(\int_X f \, d\mu_1, \ldots, \int_X f \, d\mu_n \right) \geq^{\mathcal{L}} \left(\int_X g \, d\mu_1, \ldots, \int_X g \, d\mu_n \right). \end{aligned}$$

A preference relation $\succsim$ over $\mathcal{A}(X)$ is called a ***lexicographic expected utility (LEU) preference relation*** *if there exists some $\sigma \in \mathcal{L}(X)$ such that $\succsim = \succsim_\sigma$.*

Definition 9. *Let $\sigma, \rho \in \mathcal{L}(X)$. We say that σ and ρ* ***represent the same preferences*** *and write $\sigma \cong \rho$ if $\succsim_\sigma = \succsim_\rho$. Note that $\cong$ is an equivalence relation on $\mathcal{L}(X)$.*

Definition 10. *The LPS $\sigma \in \mathcal{L}(X)$ is a* ***minimal-length*** *representation if it is the shortest LPS that represents $\succsim_\sigma$, i.e.,*

$$\forall \rho \in \mathcal{L}(X), \quad \sigma \cong \rho \implies \#\rho \leq \#\sigma.$$

For each $m \in \mathbb{N}$, the set of all minimal-length representations in $\mathcal{L}_m(X)$ is denoted by $\underline{\mathcal{L}}_m(X)$. We let $\underline{\mathcal{L}}(X) \equiv \bigcup \{\underline{\mathcal{L}}_m(X) : m \in \mathbb{N}\}$. It is obvious that every LPS has a minimal-length representation. Furthermore, the set $\underline{\mathcal{L}}_m(X)$ is nonempty for all $m \in \mathbb{N}$.

[1] Details here.

2.3 THE MAIN PROBLEM

Our main problem arises from the fact that more than one LPS can represent the same LEU preference relation.

Example 1. *Let $\sigma = (\mu_1, \mu_2) \in \mathcal{L}(X)$ and $\rho = (\mu_1, 0.5\mu_1 + 0.5\mu_2)$. It is clear that $\sigma \cong \rho$.*

Given that $\mathcal{L}(X)$ contains many redundant representations, it is reasonable to ask if there is a standard Borel space of all LEU preferences that does not contain any redundant representations. Formally, this is equivalent to asking if there is a Borel set $U \subseteq \mathcal{L}(X)$ such that the following sentence holds.

$$\forall \sigma \in \mathcal{L}(X), \exists! \rho \in U, \quad \sigma \cong \rho.$$

Furthermore, if such a U does exist then it would be useful to have a well-behaved function that maps each LPS to its equivalent representation in U. Formally, what we desire is a surjective Borel measurable map $\lambda : \mathcal{L}(X) \to \mathcal{L}(X)$ such that

$$U = \{\sigma \in U : \lambda(\sigma) = \sigma\} \text{ and } \forall \sigma, \rho \in \mathcal{L}(X), \quad \sigma \cong \rho \implies \lambda(\sigma) = \lambda(\rho).$$

It would be icing on the cake if U only contained minimal-length representations, i.e., $U \subseteq \underline{\mathcal{L}}(X)$. There are some obvious candidate techniques when X is finite or countable. However, such methods are inconceivable in the important case when X is uncountable.

The principal motivation for our problem is the construction higher-order lexicographic preferences without redundancies. Consider a finite space Θ of fundamental uncertainty. We may be interested in finding the cross section of $\mathcal{L}(X)$, where X is the uncountable set $\Theta \times \mathcal{L}(\Theta)$.

2.4 THE PROBLEM REFRAMED AND A SKETCH OF THE SOLUTION

Fortunately, we can take advantage of existing results from the field of descriptive set theory to tackle our main problem. We will reframe the question as an instance of the **Borel (cross) section problem**, which has been the subject of extensive research in mathematics. To do so, we first state two key definitions.

Definition 11. *Let $\mathbf{\Pi}$ be a partition of X. A **cross section** of $\mathbf{\Pi}$ is a subset $S \subseteq X$ such that $S \cap A$ is a singleton for every $A \in \mathbf{\Pi}$.*

Definition 12. *Let $\mathbf{\Pi}$ be a partition of X. A **section** of $\mathbf{\Pi}$ is a map $f : X \to X$ such that*

$$\forall x, y \in X, \quad x \mathbf{\Pi} f(x) \wedge (x \mathbf{\Pi} y \implies f(x) = f(y)).$$

Note that each section f defines a canonical cross section $\{x \in X : x = f(x)\}$.

It is clear that the equivalence relation $\cong$ induces a partition $\mathbf{\Pi}_{\cong}$ of $\mathcal{L}(X)$. Our problem can now be restated in the following form.

(i) Does there exist a Borel cross section of $\mathbf{\Pi}_{\cong}$?

(ii) Does there exist a Borel measurable section of $\mathbf{\Pi}_{\cong}$?

A class of results known as **Borel cross section theorems** states the conditions under which such objects exist. In order to formulate the particular cross section theorem that we use in this paper, a few more definitions are needed.

Definition 13. *Let $\mathbf{\Pi}$ be a partition of X and $A \subseteq X$. $A^* = \bigcup\{P \in \mathbf{\Pi} : A \cap P \neq \emptyset\}$ is called the **saturation** of A.*

Definition 14. *Let X be a topological space. A set $E \subseteq X$ is a G_δ set if there exists a countable family $\{G_n : n \in \mathbb{N}\}$ of open sets such that $E = \bigcap_{n\in\mathbb{N}} G_n$.*

Definition 15. *Let X be a topological space. A set $E \subseteq X$ is a F_σ set if there exists a countable family $\{F_n : n \in \mathbb{N}\}$ of closed sets such that $E = \bigcup_{n\in\mathbb{N}} F_n$.*

Informally speaking, we can think of G_δ sets and F_σ sets as almost open sets and almost closed sets, respectively. We can now state the following cross section theorem and our main result.

Proposition 1 (Theorem 5.9.1 in Srivastava [1998]). *Every partition $\mathbf{\Pi}$ of a Polish space X into G_δ sets such that the saturation of every basic open set is simultaneously F_σ and G_δ admits a section $s : X \to X$ that is Borel measurable of class 2 (i.e., the inverse image of every open set is G_δ). In particular, such partitions admit a G_δ cross section.*

Theorem 1 (The Main Theorem). *Let X be a Polish space. Then there exists a Borel set $U \subseteq \mathcal{L}(X)$ and a class-2 Borel map $\lambda\colon \mathcal{L}(X) \to \mathcal{L}(X)$ such that the following statements hold.*

(i) $U \subseteq \underline{\mathcal{L}}(X)$ *and* $\forall \sigma \in \mathcal{L}(X), \exists! \rho \in U, \quad \sigma \cong \rho$

(ii) $U = \{\sigma \in \mathcal{L}(X) : \lambda(\sigma) = \sigma\}$ *and* $\forall \sigma, \rho \in \mathcal{L}(X), \quad \sigma \cong \rho \implies \lambda(\sigma) = \lambda(\rho)$.

Furthermore, U is a Polish space if given the subspace topology with respect to $\mathcal{L}(X)$.

Theorem 1 says that there is Polish subspace U of $\mathcal{L}(X)$ that is well-behaved in several ways. First, each LPS in $\mathcal{L}(X)$ is $\cong$-equivalent to some element of U. This implies that every LEU preference relation is represented in U. Second, U does not contain redundant representations of the same LEU preference relation. Third, U is the set of all fixed points of a well-behaved function that maps each LPS to its $\cong$-equivalent representation in U.

3 PROOF OF THE MAIN THEOREM

3.1 TWO SUB-PROBLEMS

Theorem 1 follows from the two intermediate results that are stated below.

Lemma 1. *Let X be a Polish space and $m \in \mathbb{N}$. Then there exists a Borel set $U_m \subseteq \underline{\mathcal{L}}_m(X)$ and a class-2 Borel map $\varphi_m \colon \underline{\mathcal{L}}_m(X) \to \underline{\mathcal{L}}_m(X)$ such that the following statements hold.*

(i) $\forall \sigma \in \underline{\mathcal{L}}_m(X), \exists! \rho \in U_m, \quad \sigma \cong \rho$

(ii) $U_m = \{\sigma \in \underline{\mathcal{L}}_m(X) : \varphi_m(\sigma) = \sigma\}$ *and* $\forall \sigma, \rho \in \underline{\mathcal{L}}_m(X), \quad \sigma \cong \rho \implies \varphi_m(\sigma) = \varphi_m(\rho)$.

Furthermore, U_m is a Polish space if given the subspace topology with respect to $\underline{\mathcal{L}}_m(X)$.

Lemma 2. *Let X be a Polish space and $m \in \mathbb{N}$. Then there exists a class-2 Borel map $\psi_m \colon \mathcal{L}_m(X) \to \bigcup_{k \le m} \underline{\mathcal{L}}_k(X)$ such that* $\mathrm{range}(\psi_m) = \bigcup_{k \le m} \underline{\mathcal{L}}_k(X)$ *and $\forall \sigma \in \mathcal{L}_m(X), \quad \sigma \cong \psi_m(\sigma)$.*

Proof of Theorem 1. For each $m \in \mathbb{N}$, let U_m, φ_m, and ψ_m be the objects that exist by Lemmas 1 and 2. Let $U \equiv \bigcup_{m \in \mathbb{N}} U_m$ and define λ as follows.

$$\forall \sigma \in \mathcal{L}(X), \quad \lambda(\sigma) \equiv \varphi_{\#\psi(\sigma)}(\psi_{\#\sigma}(\sigma))$$

The desired properties of U and λ follow immediately. □

Sections 3.2 and 3.3 cover the proofs of Lemmas 2 and 1, respectively.

3.2 THE ANATOMY OF A LEXICOGRAPHIC PROBABILITY SYSTEM

The proof of Lemma 2 can be more easily understood by deconstructing the anatomy of each LPS σ as it relates to the determination of the LEU preference relation that is represented by σ.

Definition 16. *Let $m \in \mathbb{N}$ and $\sigma = (\mu_1, \dots, \mu_m) \in \mathcal{L}(X)$. For all $k \in \{1, \dots, \#\sigma\}$, let $\sigma|k$ denote the **length-k initial segment** of σ.*

$$\sigma|k \equiv (\mu_1, \dots, \mu_k)$$

Definition 17. *Let $m \in \mathbb{N}$ and $\sigma = (\mu_1, \dots, \mu_m) \in \mathcal{L}(X)$. The **relevent index set** of σ, denoted by $\mathbb{I}_\sigma$, is defined as follows. Members of $\mathbb{I}_\sigma$ are called **relevant indices** of σ.*

$$\mathbb{I}_\sigma \equiv \{k \in \mathbb{N} : k \le m \wedge \forall j < k, \quad (\sigma|k \ncong \sigma|j)\}$$

The notion of a relevant index is an intuitive one. Suppose that $\sigma = (\mu_1, \ldots, \mu_m) \in \mathcal{L}(X)$ is an LPS and $1 \leq k \leq m$. If there exists a $j < k$ such that $\sigma|k \cong \sigma|j$, then μ_k adds no information about preferences that is not already captured by $\sigma|j$. If so, then omitting μ_k from the description of σ would not remove any information that is *relevant* to the evaluation of preferences. It should be immediately apparent that the minimal-length representation of $\succsim_\sigma$ must have length equal to the size of the set $\mathbb{I}_\sigma$. A minimal-length representation can be constructed by simply stripping σ of all components that have irrelevant indices.

Example 2. *Let $\mu, \nu \in \mathcal{P}(X)$ and $\mu \neq \nu$.*

$$\begin{aligned}
\sigma &= (\mu_1, \mu_2, \mu_3, \mu_4) = (\nu, \nu, 0.8\nu + 0.2\mu, \mu) \in \mathcal{L}(X) \\
\mathbb{I}_\sigma &= \{1, 3\} \\
\sigma &\cong (\mu_1, \mu_3) \in \underline{\mathcal{L}}(X)
\end{aligned}$$

Proof of Lemma 2. For each $D \in \mathbf{2}^{\{1,\ldots,m\}} \setminus \{\emptyset\}$, let $A_D \equiv \{\sigma \in \mathcal{L}_m(X) : \mathbb{I}_\sigma = D\}$. Therefore, $A = \{A_D : D \in \mathbf{2}^{\{1,\ldots,m\}} \setminus \{\emptyset\}\}$ is a finite partition of $\mathcal{L}_m(X)$ into nonempty closed sets. We can define ψ_m in a piecewise fashion on each $A_D \in A$ by letting $\psi_m(\sigma) \equiv (\mu_k)_{k \in D}$ for each $\sigma = (\mu_1, \ldots, \mu_m) \in A_D$. Therefore, the map ψ_m is piecewise continuous over each $A_D \in A$. The inverse image of an open set under ψ_m is therefore a finite union of sets that are each finite intersections of open sets with closed sets. Therefore, the inverse image of an open set under ψ_m is a G_δ set. It follows that ψ_m is a class-2 Borel map.[2] $\square$

3.3 PROOF OF LEMMA 1

In order to prove Lemma 1, we must show that the partition $\mathbf{\Pi}_\cong$ of the Polish space $\underline{\mathcal{L}}_m(X)$ satisfies the premises for the application of Proposition 1. In other words, we must show that

(i) each $P \in \mathbf{\Pi}_\cong$ is G_δ; and

(ii) the saturation of every basic open set is simultaneously F_σ and G_δ.

Remark 1. *$\mathcal{P}(X)$ is a closed convex subset of a locally convex Hausdorff topological vector space, namely the space $\mathcal{M}(X)$ of totally finite and countably additive signed borel measures on X with the topology of weak convergence. $\mathcal{P}(X)$ is a topological subspace of $\mathcal{M}(X)$, but not its vector subspace.*

Remark 2. *The set $\underline{\mathcal{L}}_m(X)$ is open in the Polish space $\mathcal{L}_m(X)$ for all $m \in \mathbb{N}$.*

Lemma 3 (Part i)**.** *Let $m \in \mathbb{N}$ and $\sigma = (\mu_1, \ldots, \mu_m) \in \underline{\mathcal{L}}_m(X)$. The $\cong$-equivalence class $\mathbf{\Pi}_\sigma$ of σ is a G_δ set in $\underline{\mathcal{L}}_m(X)$*

[2] It can also be shown that ψ_m is not a class-1 Borel map (i.e., ψ_m is not continuous).

Proof of Lemma 3. Note that $\mathbf{\Pi}_\sigma$ is the following product set.

$$\mathbf{\Pi}_\sigma \equiv \prod_{k=1}^{m} \{\nu_k \in \mathcal{P}(X) : \exists(\alpha_j) \in \mathbb{R}^k \quad \textstyle\sum_{j=1}^k \alpha_j = 1 \wedge \alpha_k > 0 \wedge \nu_k = \sum_{j=1}^k \alpha_j \mu_j\}$$

We only need to prove that the following is G_δ for all $k \leq m$.

$$E_k \equiv \{\nu_k \in \mathcal{P}(X) : \exists(\alpha_j) \in \mathbb{R}^k \quad \textstyle\sum_{j=1}^k \alpha_j = 1 \wedge \alpha_k > 0 \wedge \nu_k = \sum_{j=1}^k \alpha_j \mu_j\}$$

Given that $\sigma \in \underline{\mathcal{L}}_M(X)$, the measures $\mu_1, \ldots, \mu_m$ are linearly independent. It follows that the closed set

$$\operatorname{span}(\sigma|k) \equiv \{\nu_k \in \mathcal{M}(X) : \exists(\alpha_j) \in \mathbb{R}^k \quad \nu_k = \textstyle\sum_{j=1}^k \alpha_j \mu_j\}$$

of signed measures spanned by $\mu_1, \ldots, \mu_k$ is homeomorphic to $\mathbb{R}^k$ and the map $(\alpha_j) \mapsto \sum_{j=1}^k \alpha_j \mu_j$ is a homeomorphism between the two spaces. It follows that $E_k = E_k' \cap \mathcal{P}(X)$, where E_k' is defined as follows.

$$E_k' \equiv \{\nu_k \in \mathcal{M}(X) : \exists(\alpha_j) \in \mathbb{R}^k \quad \textstyle\sum_{j=1}^k \alpha_j = 1 \wedge \alpha_k > 0 \wedge \nu_k = \sum_{j=1}^k \alpha_j \mu_j\}$$

It is immediate that E_k' is homeomorphic to a G_δ subset of $\mathbb{R}^k$. Therefore, E_k' is also G_δ in $\mathcal{M}(X)$ because it is G_δ in the closed topological subspace $\operatorname{span}(\sigma|k)$. Then E_k is the intersection of the G_δ set E_k' with $\mathcal{P}(X)$. Therefore, E_k is a G_δ in $\mathcal{P}(X)$. □

Definition 18. *Let T be a topological vector space and let $V \subseteq T$. A set $U \subseteq V$ is* ***radial*** *in V if there is some $u \in U$ such that, for each $v \in V$,*

(i) $\exists u_v \in U$ such that U includes the line segment joining u and u_v; and

(ii) v lies on the ray that originates at u and passes through u_v.[3]

Remark 3. *Let V be a locally convex topological vector space. Let C be a closed convex subset of V. Let U be a radial convex open set such that $U \cap C \neq \emptyset$. Let $v \in V$. Then the following set is open in the topological subspace C.*

$$\{z \in C : z \in \alpha(U \cap C) + (1-\alpha)v \wedge \alpha > 0\}$$

Lemma 4 (Part ii)**.** *Let $m \in \mathbb{N}$ and $\sigma = (\mu_1, \ldots, \mu_m) \in \underline{\mathcal{L}}_m(X)$. Let U be a basic open set in $\underline{\mathcal{L}}_m(X)$. The saturation U^*, defined below, is open.*

$$U^* \equiv \bigcup_{\sigma \in U} \mathbf{\Pi}_\sigma$$

[3] This definition differs slightly from the more common one found in functional analysis texts such as Aliprantis and Border [1999, p. 168]. It is more usual to fix $u = 0 \in T$ and define radial neighborhoods of zero in T. Definition 18 says that a set U is radial in V if it is a nonempty intersection of V with with a radial neighborhood of $0 \in T$.

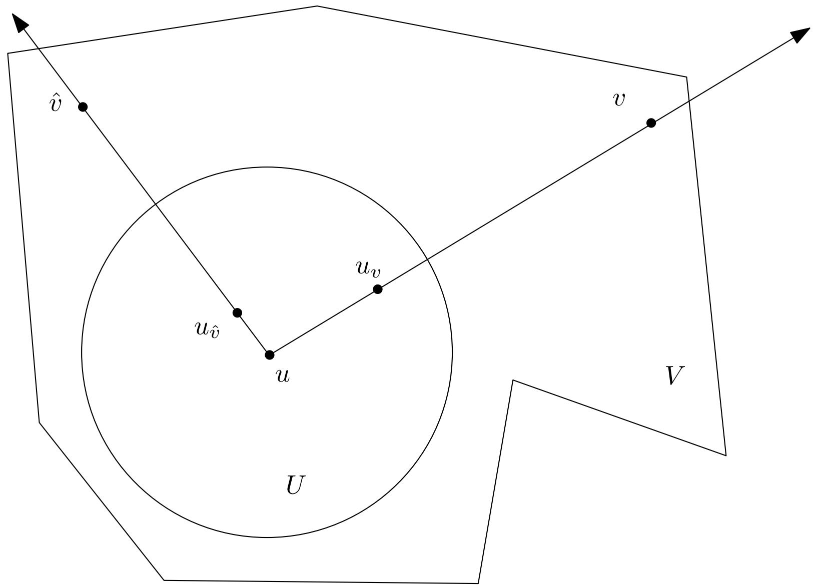

Figure 1: An example of U radial in V. Each $v \in V$ is reached by a ray that originates at u and contains a line segment in U that has u as one of its endpoints.

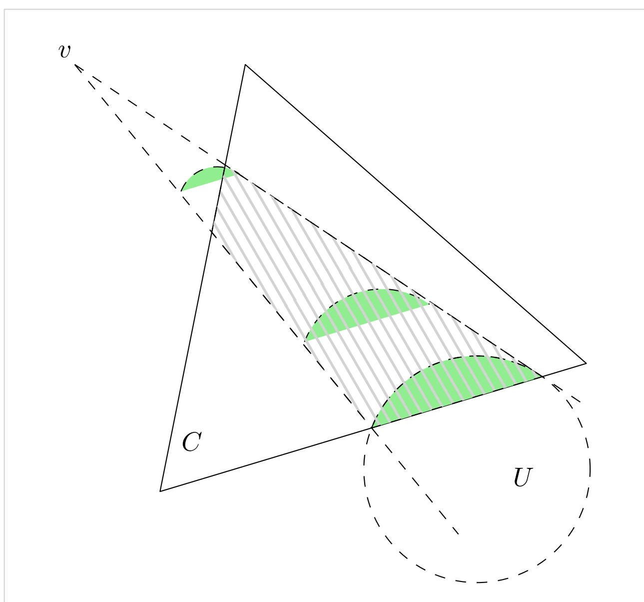

Figure 2: An example of $\alpha(U \cap C) + (1-\alpha)v$ in Remark 3

Proof of Lemma 4. The proof is by induction on m. The base case ($m = 1$) is immediate because $\mathcal{P}(X) = \mathcal{L}_1(X) = \underline{\mathcal{L}}_1(X)$.

Fix an $m \in \mathbb{N}$. Suppose that the saturation of any basic open set in $\underline{\mathcal{L}}_m(X)$ is open. Any basic open set U in $\underline{\mathcal{L}}_{m+1}(X)$ can be written as $(V \times O) \cap \underline{\mathcal{L}}_{m+1}(X)$, where

V is a basic open set in $\underline{\mathcal{L}}_m(X)$ and O a basic open set in $\mathcal{P}(X)$. Then we can write U^* as follows.

$$\begin{aligned} U^* &= \{\sigma.(\mu) \in \underline{\mathcal{L}}_m(X) : \sigma \in V^* \wedge \mu \in \alpha O + (1-\alpha)\operatorname{span}(\sigma) \wedge \alpha > 0\} \\ &= \bigcup_{\sigma \in V^*} \{\sigma.(\mu) \in \underline{\mathcal{L}}_m(X) : \mu \in \alpha O + (1-\alpha)\operatorname{span}(\sigma) \wedge \alpha > 0\} \end{aligned}$$

V^* is open by the induction hypothesis. We need only prove that the following set is open in $\mathcal{P}(X)$.[4]

$$\begin{aligned} &\{\mu \in \mathcal{P}(X) : \mu \in \alpha O + (1-\alpha)\operatorname{span}(\sigma) \wedge \alpha > 0\} \\ = &\bigcup_{\rho \in \operatorname{span}(\sigma)} \{\mu \in \mathcal{P}(X) : \mu \in \alpha O + (1-\alpha)\rho \wedge \alpha > 0\} \end{aligned}$$

In light of Remark 1, we can assume that O is convex and radial without loss of generality because locally convex topological vector spaces have neighborhood base systems of convex radial open sets. The subspace topology of $\mathcal{P}(X)$ has neighborhood base systems that inherit those properties because $\mathcal{P}(X)$ is a closed convex subset of $\mathcal{M}(X)$. Remark 3 implies that the set defined above is a union of open sets and therefore itself open. □

REFERENCES

Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis*. Springer-Verlag, New York, 1999.

Lawrence Blume, Adam Brandenburger, and Eddie Dekel. Lexicographic probabilities and choice under uncertainty. 59(1):61–79, January 1991. URL http://www.jstor.org/stable/2938240.

Adam Brandenburger, Amanda Friedenberg, and H. Jerome Keisler. Admissibility in games. 76(2):307–352, March 2008. URL http://www.jstor.org/stable/40056426.

Alfré Rényi. On a new axiomatic theory of probability. *Acta Mathematica Hungarica*, 6(3):285–335, November 1955. URL http://dx.doi.org/10.1007/BF02024393.

S. M. Srivastava. *A Course on Borel Sets*. Springer-Verlag, New York, 1998.

[4] In a product space $X \times Y$, any subset of the form $E = \bigcup_{x \in U}\{x\} \times V_x$ is open when $U \subseteq X$ is open and $V_x \subseteq Y$ is open for all $x \in U$. Note that $E^c = (U^c \times Y) \cup (\bigcup_{x \in U}\{x\} \times V_x^c)$ is closed under those premises.