Confidence Models of Incomplete Preferences*

Morgan McClellon[†]

May 26, 2015

Abstract

This paper introduces and axiomatizes a new class of representations for incomplete preferences called confidence models. Confidence models describe decision makers who behave as if they have probabilistic uncertainty over their true preferences, and are only willing to express a binary preference if it is sufficiently likely to hold. Confidence models are flexible enough to model behavior on a variety of domains, and they are general enough to nest the popular multi-utility models of incomplete preferences. Most importantly, they provide a natural way to connect incomplete preferences with stochastic choice. This connection is characterized by a simple condition that serves to identify the behavioral content of incomplete preferences.

Keywords: Confidence models; incomplete preferences; random choice.

*I would like to thank Tomasz Strzalecki for excellent advice and support throughout this project. I also benefited from helpful conversations with Mira Frick, Drew Fudenberg, and Jerry Green, and feedback from two anonymous referees. All remaining errors are my own. This paper is based on a chapter of my dissertation at Harvard University and was partially funded by a Graduate Research Fellowship from the National Science Foundation.

[†]Contact mccllellon.morgan@gmail.com.

1 Introduction

This paper introduces and axiomatizes a new class of representations for incomplete preferences called confidence models.¹ A decision maker (DM) with preferences represented by a confidence model has probabilistic uncertainty over her underlying preferences. When comparing two alternatives, the DM computes the probability that one is weakly preferred to the other, and expresses a weak preference if and only if this probability exceeds a fixed threshold. Intuitively, the DM only expresses binary preferences in which she holds sufficient confidence.

Confidence models can be formulated for incomplete preferences on a variety of domains. Here, we define and axiomatize confidence models for preferences over a finite set and for preferences over lotteries. We show that confidence models are remarkably general and nest the prominent existing models of incomplete preferences on both these domains. Potential extensions to additional domains are discussed in the conclusion.

The generality of confidence models allows them to encompass behavior outside the scope of existing models for incomplete preferences. In particular, preferences represented by confidence models need not be transitive. When interpreting incomplete preferences in the context of political science, failures of transitivity arise naturally due to Condorcet cycles. Confidence preferences can accommodate such phenomena and can be interpreted as a model of supermajority aggregation rules. The potential lack of transitivity distinguishes confidence models from the popular multi-utility models of incomplete preferences axiomatized by Evren and Ok (2008) and Dubra et al. (2004).

Confidence models suggest a natural way to connect incomplete preferences and stochastic choice. Specifically, an incomplete preference and a random choice rule are defined to be consistent if they can be represented using a common probability measure on the space of complete preference relations. We show that this notion of consistency has a simple characterization in terms of a joint condition on the incomplete preferences and the stochastic choice data. The tight connection between confidence preferences and random choice provides a novel way to understand the behavioral content of incomplete preferences.

In a recent experiment, Agranov and Ortoleva (2015) found that subjects chose stochastically even when the same menu was presented to them multiple times in a row. Afterwards, subjects were asked to explain their choices, and many indicated that they chose stochastically as a result of being unable to select an optimal alternative. To be sure, this is anecdotal evidence, but it certainly suggests that stochastic choice may arise at the individual level as

¹It is worth mentioning that the term “confidence model” has been used in a different decision-theoretic context by Chateauneuf and Faro (2009). That paper studies complete preferences over Anscombe-Aumann acts and characterizes a representation that features a fuzzy set of priors.

a result of incomplete preferences.

Section 2 studies confidence preferences on a finite set of alternatives, while Section 3 studies preferences on lotteries over a finite set of alternatives. Limitations on the uniqueness of confidence representations are discussed in Section 4. Section 5 discusses related literature, and in particular formalizes the relationship between confidence models and multi-utility models.

2 Confidence Models on Finite Sets of Alternatives

2.1 Basic Definitions

Let X be a set of alternatives with $|X| = N$ and $2 \leq N < \infty$. Let $\mathcal{W}$ denote the set of all complete and transitive relations on X . A binary relation R is antisymmetric if xRy and yRx implies $x = y$; let $\mathcal{R} \subseteq \mathcal{W}$ denote the set of all complete, transitive, antisymmetric relations on X . A traditional agent with complete preferences is described by a preference relation $W \in \mathcal{W}$; if the agent is allowed to express only strict preferences between distinct objects, she is described by a preference relation $R \in \mathcal{R}$.

We are interested in a DM with a potentially-incomplete preference relation $\succsim$ on X , who behaves as if she is unsure of her “true” preferences and is willing to express only binary preferences in which she is sufficiently confident. Formally, such preferences are represented as follows:

Definition 1. A *confidence model* for a preference $\succsim$ is a pair (μ, α) where $\mu \in \Delta(\mathcal{W})$ and $\alpha \in (1/2, 1]$ such that for any $x, y \in X$ we have

$$x \succsim y \quad \text{if and only if} \quad \mu(\{W \mid xWy\}) \geq \alpha.$$

A *regular confidence model* for $\succsim$ is a confidence model (μ, α) where $\text{supp}(\mu) \subseteq \mathcal{R}$.

In a confidence model, the measure μ captures the DM’s uncertainty over her preferences, and the parameter α captures the degree of confidence the DM must have before expressing a given binary preference. Specifically, when comparing two alternatives x and y , the DM behaves as if she is computing the μ -probability that x is preferred to y and expressing the preference $x \succsim y$ if and only this probability exceeds the threshold α .

A DM with preferences captured by a regular confidence model behaves as if she is certain that her preferences will not contain any nontrivial indifference. The distinction between confidence models and regular confidence models will be important in Section 2.3 when we connect confidence preferences to random choice.

2.2 Axiomatization

The following proposition shows that confidence models are extremely general: they can represent any reflexive preference. The subclass of regular confidence models is characterized by antisymmetry.

Corollary 1. *Let $\succsim$ be a binary relation on X .*

1. *$\succsim$ is reflexive if and only if it can be represented by a confidence model.*
2. *$\succsim$ is reflexive and antisymmetric if and only if it can be represented by a regular confidence model.*

This result is a corollary of McGarvey’s Theorem, a classic result in social choice theory established by McGarvey (1953).² McGarvey’s Theorem states that any complete relation can arise as the outcome of majority rule applied a collection of complete, transitive and antisymmetric preferences.

Corollary 1 can be established directly from McGarvey’s Theorem with straightforward arguments. The appendix provides a novel proof for Corollary 1 that is not based on McGarvey’s Theorem, but instead relies on the connection between confidence preferences and stochastic choice that is formalized in Lemma 1. An outline of this proof may be found after the statement of Lemma 1 in Section 2.3. The (non)uniqueness of confidence representations is discussed in Section 4.

The following example illustrates how confidence preferences may fail to be transitive when the measure μ supports a Condorcet cycle.

Example 1. Let $X = \{x, y, z\}$. Consider the following relations from $\mathcal{R}$:

$$R_1: xR_1yR_1z;$$

$$R_2: zR_2xR_2y;$$

$$R_3: yR_3zR_3x.$$

Let μ assign probability 1/3 to each relation R_1 , R_2 , and R_3 . Let $\alpha = .6$. If $\succsim$ is represented by (μ, α) then it is straightforward to check that $x \succsim y$ and $y \succsim z$, but $\neg(x \succsim z)$, in violation of transitivity. $\diamond$

Consistent with the connection to McGarvey’s theorem, confidence preferences can be given a social-choice interpretation. In this view, $\succsim$ represents the preferences of “Society,”

²I thank an anonymous referee for pointing out this connection.

which are formed by aggregating the (complete and transitive) preferences of individuals. If $\succsim$ has a confidence representation, we can interpret μ as the distribution of preferences in the population, and α as the threshold for a supermajority voting rule.

2.3 Regular Confidence Preferences and Stochastic Choice

Confidence models provide a natural way to connect incomplete preferences with random choice data. Before formalizing this connection, we briefly review the essential concepts from the literature on stochastic choice.

Let X again be a finite set of alternatives with $|X| = N$, and let $\mathcal{D}$ denote the set of all nonempty subsets of X . A *random choice rule (RCR)* is a function $\rho: \mathcal{D} \times X \rightarrow [0, 1]$, where $\rho^D(x) := \rho(D, x)$ is interpreted as the observed probability that x is chosen from D . For this to make sense, we impose that for all $D \in \mathcal{D}$

$$\sum_{x \in D} \rho^D(x) = 1.$$

In the context of random choice, a probability measure $\mu \in \Delta(\mathcal{R})$ is called a *random utility function (RUF)*. A RUF is said to *represent the RCR* ρ if

$$\rho^D(x) = \mu(\{R \in \mathcal{R} \mid x \text{ is } R\text{-optimal in } D\})$$

for all $(D, x) \in \mathcal{D} \times X$. A classic result established by Falmagne (1978) and rediscovered by Barberá and Pattanaik (1986) shows that a RCR ρ can be represented by some RUF if and only if the following inequalities hold:

$$\sum_{C: D \subseteq C} (-1)^{|C \setminus D|} \rho^C(x) \geq 0 \quad \forall (D, x) \in \mathcal{D} \times X.$$

These inequalities are known as the *Block-Marschak (BM) inequalities*;³ they are written here in the form of Fiorini (2004).

In the language of stochastic choice, regular confidence models represent a reflexive, antisymmetric preference relation using two components: a RUF and a threshold. This suggests a natural connection between such preferences and stochastic choice.

Definition 2. A reflexive, antisymmetric preference $\succsim$ is *consistent* with a RCR ρ satisfying the BM inequalities if there exists a RUF $\mu \in \Delta(\mathcal{R})$ such that: (1) (μ, α) represents $\succsim$ for

³In honor of Block and Marschak (1960), which established that they are necessary for a RCR to be represented by a RUF.

some $\alpha \in (1/2, 1]$; and (2) μ represents ρ .

In words, the preference $\succsim$ is consistent with stochastic choice data if we can find representations for each that share a common RUF. Although we have defined consistency in terms of the representations for $\succsim$ and ρ , there is a natural condition that allows an observer to check directly whether an incomplete preference is consistent with a given set of stochastic choice data.

Definition 3 (Binary Compatibility). A preference $\succsim$ and a RCR ρ are *compatible on binary sets* if for all $x \succsim y$ and $\neg(x' \succsim y')$, we have $\rho^{\{x,y\}}(x) > \rho^{\{x',y'\}}(x')$.

Binary compatibility requires that if x is weakly preferred to y while x' is not weakly preferred to y' (either because $y' \succ x'$ or because x' and y' are not comparable) then x must be chosen with a higher frequency from the binary menu $\{x, y\}$ than x' is chosen from the menu $\{x', y'\}$.

The following lemma shows that binary compatibility characterizes consistency between preferences and stochastic choice.

Lemma 1. *Let $\succsim$ be a reflexive, antisymmetric preference and let ρ be a RCR satisfying the BM inequalities. Then $\succsim$ and ρ are consistent if and only if they are compatible on binary sets.*

Proof. Assume that $\succsim$ and ρ are compatible on binary sets. Let μ be a RUF representing ρ . Let $S = \{(x, y) \mid x \succsim y\}$ and define

$$\alpha := \min_S \rho^{\{x,y\}}(x).$$

Because $\succsim$ is reflexive, S is nonempty and $\alpha \leq 1$. Together, the antisymmetry of $\succsim$ and compatibility of $\succsim$ and ρ on binary sets further imply that $\alpha > 1/2$. Recalling that μ represents ρ , binary compatibility and the definition of α imply that $\mu\{R \in \mathcal{R} \mid x R y\} \geq \alpha$ if and only if $x \succsim y$, showing that (μ, α) is regular confidence model for $\succsim$. This representation features the same RUF used to represent ρ , which establishes consistency between $\succsim$ and ρ .

Conversely, assume $\succsim$ and ρ are consistent—that is, there exists a $\mu \in \Delta(\mathcal{R})$ such that μ represents ρ and (μ, α) represents $\succsim$ for some $\alpha \in (1/2, 1]$. Let $x \succsim y$ and $\neg(x' \succsim y')$. Then

$$\rho^{\{x,y\}}(x) = \mu(\{R \in \mathcal{R} \mid x R y\}) \geq \alpha > \mu(\{R \in \mathcal{R} \mid x' R y'\}) = \rho^{\{x',y'\}}(x'),$$

showing that $\succsim$ and ρ are compatible on binary sets. In the preceding equation, the equalities hold because μ represents ρ and the inequalities hold because (μ, α) represents $\succsim$. ■

Lemma 1 provides a tractable characterization of consistency between reflexive, antisymmetric preferences and stochastic choice data. In addition, it provides a useful technical tool that plays a key role in the proof of Theorem 1. Indeed, the proof of Theorem 1(B) takes a reflexive, antisymmetric preference $\succsim$ and constructs a hypothetical RCR that satisfies the BM inequalities. This RCR is constructed to be compatible with $\succsim$ on binary sets, allowing Lemma 1 to deliver a confidence model for $\succsim$.

The proof of Theorem 1(A) also uses Lemma 1 to construct a measure $\mu \in \Delta(\mathcal{R})$. However, the hypothetical RCR used to generate this measure is not necessarily compatible with $\succsim$ on binary sets, and the resulting RUF must be modified before it can be used in a confidence representation of $\succsim$. Specifically, μ is mixed with several different measures on $\mathcal{W}$ to form a new measure $\mu' \in \Delta(\mathcal{W})$. This measure is then combined with an appropriate threshold to generate a confidence model for $\succsim$. Details may be found in the appendix.

3 Confidence Models on Sets of Lotteries

3.1 Setup and Axiomatization

The intuition underlying confidence models—a DM is unsure of her preferences, and expresses only those she believes will hold with sufficiently high probability—does not depend on the structure of the domain X . In principle, then, confidence models can be defined on a variety of domains. In this section we study confidence models for incomplete preferences over lotteries. Several technical issues arise when characterizing confidence models, but the strong connection to stochastic choice remains.

Let X again be a finite set with $|X| = N$. The primitive for this section is a preference relation $\succsim$ defined on a set of $\mathcal{L} \subseteq \Delta(X)$. Endow the domain $\mathcal{L}$ with the standard topology it inherits as a subset of $\mathbb{R}^N$.

The set of twice normalized expected utility functions on $\Delta(X)$ is

$$\mathcal{U} := \left\{ u \in \mathbb{R}^n \mid \sum_i u_i = 0, \sum_i u_i^2 = 1 \right\}.$$

Endow $\mathcal{U}$ with the standard Euclidean topology, and let $\Delta(\mathcal{U})$ be the set of finitely-additive measures on the Borel σ -algebra of $\mathcal{U}$.

Confidence models for incomplete preferences over lotteries are defined as follows.

Definition 4. A *confidence model* for a preference $\succsim$ over $\mathcal{L}$ is a pair (μ, α) where $\mu \in \Delta(\mathcal{U})$

and $\alpha \in (1/2, 1]$, such that for all $p, q \in \mathcal{L}$

$$p \succsim q \quad \text{iff} \quad \mu(\{u \in \mathcal{U} \mid u \cdot p \geq u \cdot q\}) \geq \alpha.$$

Confidence models describe DMs who are unsure of their underlying expected utility preferences, and express a preference for p over q if and only if their subjective measure μ assigns a sufficiently high probability to the expected utility of p exceeding the expected utility of q .

In Section 2, we defined regular confidence models, which featured measures that put zero probability on the DM being indifferent between two alternatives. Regular confidence models for preferences over lotteries are defined analogously: they use measures on $\mathcal{U}$ for which ties are a zero-probability event.

Definition 5. A measure $\mu \in \Delta(\mathcal{U})$ is *regular* if any finite collection of lotteries has a unique maximal element with μ -probability 1.⁴ A confidence model (μ, α) featuring a regular μ is called a *regular confidence model*.

The main theorem of this section characterizes confidence models when $\mathcal{L}$ is finite.

Theorem 1. Let $\mathcal{L} \subseteq \Delta(X)$ be finite. A preference $\succsim$ on $\mathcal{L}$ can be represented by a confidence model if and only if it is reflexive and satisfies Independence, namely: for any $p, q \in \mathcal{L}$, $r \in \Delta(X)$ and $\lambda \in (0, 1)$,

$$p \succsim q \quad \Rightarrow \quad \lambda p + (1 - \lambda)r \succsim \lambda q + (1 - \lambda)r$$

whenever both mixtures in the implication are in $\mathcal{L}$.

Compared to Theorem 1, this result adds only the standard Independence axiom from expected utility theory to characterize confidence models on finite sets of lotteries. Clearly this axiom is necessary: every $u \in \mathcal{U}$ represents a complete preference over lotteries that satisfies Independence. Because the domain $\mathcal{L}$ is finite, there is no need for a continuity axiom.

The case of a finite $\mathcal{L}$ holds some interest, as experiments necessarily can elicit only a finite number of preferences. From a theoretical viewpoint, the other natural case to consider is $\mathcal{L} = \Delta(X)$, the standard domain for axiomatizing preferences over lotteries. I conjecture, but have been unable to prove, that in this case regular confidence models are characterized by reflexivity, antisymmetry, and Independence. Surprisingly, the standard continuity axiom

⁴For a complete formal definition, see Gul and Pesendorfer (2006).

does not necessarily hold unless the measure μ is countably additive.⁵ In addition, merely dropping antisymmetry will no longer suffice to distinguish confidence models from regular confidence models. See Appendix B for details on these technical issues.

3.2 Stochastic Choice Over Lotteries

Section 2.3 showed that regular confidence models on finite sets of alternatives have a natural connection to stochastic choice. A similar connection exists for regular confidence models on lotteries. We begin by redefining notation to match the new domain. Throughout this section, we consider preferences on the domain $\mathcal{L} = \Delta(X)$.

Let $\mathcal{D}$ denote the set of all nonempty, finite subsets of $\Delta(X)$. Let $\Delta(\Delta(X))$ denote the set of all simple probability measures on $\Delta(X)$. A *random choice rule (RCR)* is function $\rho: \mathcal{D} \rightarrow \Delta(\Delta(X))$, where we let $\rho^D(p)$ denote the probability that p is chosen from D , and require that $\rho^D(D) := \sum_{p \in D} \rho^D(p) = 1$ for all D .

Gul and Pesendorfer (2006) introduce the following (paraphrased) axioms for a RCR:

GP1 Mixture continuity: $\rho^{\lambda D + (1-\lambda)D'}$ is continuous in λ .

GP2 Monotonicity: if $p \in D \subset D'$, then $\rho^{D'}(p) \leq \rho^D(p)$.

GP3 Linearity: $\rho^{\lambda D + (1-\lambda)\{q\}}(\lambda p + (1-\lambda)\{q\}) = \rho^D(p)$.

GP4 Extremeness: $\rho^D(\text{ext}(D)) = 1$.

In the context of random choice over lotteries, a measure $\mu \in \Delta(\mathcal{U})$ is called a random utility function (RUF). An RCR ρ is *maximized* by a RUF μ if $\rho^D(p)$ is equal to the μ -probability that p is optimal in D . Gul and Pesendorfer (2006) prove that an RCR satisfies GP1–4 if and only if it is maximized by a regular measure ν , which is defined on a set of expected utility functions that are normalized differently than the expected utilities in $\mathcal{U}$. These differences are addressed by Ahn and Sarver (2013), who prove that an RCR ρ satisfies Axioms GP1–4 if and only if it is maximized by a unique regular $\mu \in \Delta(\mathcal{U})$.

We are now ready to introduce the analogue of Definition 2 for preferences over lotteries. The term "regular confidence preference" refers a preference that has a regular confidence representation.

Definition 6. A regular confidence preference $\succsim$ on $\mathcal{L} = \Delta(X)$ is *consistent* with a RCR ρ satisfying GP1–4 if there exists an $\alpha \in (1/2, 1]$ such that the regular confidence model (μ, α) represents $\succsim$, where μ is the unique, regular RUF that maximizes ρ .

⁵I thank an anonymous referee for calling attention to this issue and correcting an error in a previous version of this paper.

The following lemma shows that consistency is once again characterized by binary compatibility.

Definition 7 (Binary Compatibility). A RCR ρ is *compatible on binary sets* with a preference $\succsim$ if for all pairs of lotteries (p, p') and (q, q') with $p \succsim p'$ and $\neg(q \succsim q')$, we have $\rho^{\{p, p'\}}(p) > \rho^{\{q, q'\}}(q)$.

Lemma 2. *A RCR ρ satisfying GP1–4 is compatible on binary sets with a regular confidence preference $\succsim$ that satisfies the axioms of Theorem 1 if and only if ρ and $\succsim$ are consistent.*

The proof (similar in spirit to that of Lemma 1) may be found in the appendix. If an analyst is confronted with a DM who reports an incomplete preference relation and is observed to choose stochastically, Lemma 2 allows the analyst to check whether the reported preferences and the observed data are consistent. In this sense, Lemma 2 identifies the behavioral content of incomplete preferences and shows that it rests in the frequency of choice from binary menus.

Lemma 2 is relevant for selecting a “canonical” confidence model to represent a given preference, as discussed in the next section.

4 Uniqueness

Confidence models feature two parameters and hoping for both to be uniquely identified from preferences is clearly overoptimistic. A more reasonable expectation is for “conditional uniqueness”; that is, for α to be unique given μ and vice-versa. Unfortunately, Example 2 demonstrates that for the domains studied in this paper, even conditional uniqueness is too much to ask.

Example 2. Consider a reflexive, antisymmetric preference $\succsim$ on $X = \{x, y\}$. The set $\mathcal{R}$ consists of two relations: R_1 strictly prefers good x while R_2 strictly prefers good y . Suppose a DM’s preferences are represented by the (regular) confidence model (μ, α) where

$$\mu(R_1) = 3/4, \mu(R_2) = 1/4, \text{ and } \alpha = 2/3.$$

A brief inspection reveals that this DM’s preferences are actually identical to a standard DM with preference relation R_1 . Moreover, fixing μ this will still be true for any $\alpha \in (1/2, 3/4]$. Alternatively, if we fix $\alpha = 2/3$ then preferences are the same for any $\mu(R_1) \in [2/3, 1]$.

This example can easily be modified to illustrate nonuniqueness for preferences over lotteries. After normalizing, there are two strict expected utility functions on $\Delta(X)$: u_1

strictly prefers the lottery δ_x and u_2 strictly prefers the lottery δ_y . So, if we define $\mu \in \Delta(\mathcal{U})$ by $\mu(u_1) = 3/4$ and $\mu(u_2) = 1/4$ then the argument given above shows that the preferences on $\Delta(X)$ represented by $(\mu, 2/3)$ have many different confidence representations. $\diamond$

Facing this lack of uniqueness, one might hope to identify a “canonical” representation for a given confidence preference. In light of the tight connection between confidence preferences and random choice, a natural way to approach this problem is by gathering stochastic choice data. Specifically, suppose we have preference $\succsim$ on $\Delta(X)$ with a regular confidence representation, and we observe the DM make stochastic choices captured by an RCR ρ . If $\succsim$ and ρ are found to be compatible on binary sets, then the proof of Lemma 2 provides a confidence model (μ, α) that is an excellent candidate for a canonical representation of $\succsim$. The measure μ is the unique RUF that maximizes the DM’s observed choice data, and the threshold α is easily seen to be the maximal threshold that can represent $\succsim$ in conjunction with μ .

For preferences over finite sets, no such canonical representation can be identified: it is well known that a RCR ρ on X can be represented by more than one RUF $\mu \in \Delta(\mathcal{R})$ (see Fishburn (1998) for an example and McClellon (2015) for extensive discussion of this issue). Nonetheless, the threshold α constructed in the proof of Lemma 1 is still the maximal threshold that can represent $\succsim$ in conjunction with any μ that represents ρ .

5 Related Literature

5.1 Multi-Utility Models and Transitivity

For finite X , Theorem 1(A) shows that confidence models nest any reflexive preference. This includes, of course, the “traditional” model for an economic DM: a single preference relation $W \in \mathcal{W}$ (take $\mu = \delta_W$ and any $\alpha > 1/2$). It also includes the multi-utility model axiomatized by Evren and Ok (2008). Those authors show that any reflexive and transitive preference can be represented by a set of weak orders $\mathcal{M} \subseteq \mathcal{W}$:

$$x \succsim y \quad \Leftrightarrow \quad xWy \text{ for all } W \in \mathcal{M}.$$

The interpretation of this model is that the DM considers a set of preference orderings to be possible, and requires unanimous agreement among these preferences before she is willing to express a preference. Unlike in confidence models, the DM does not express any probability judgments among the preferences in $\mathcal{M}$.

Like Multi-Utility models, Confidence models with $\alpha = 1$ require unanimous agreement

before any preference is expressed. A simple corollary of Evren and Ok (2008)'s theorem shows that the existence of such a representation is characterized by transitivity.

Corollary 2. *A preference $\succsim$ on a finite set X is reflexive and transitive if and only if it can be represented by a confidence model (μ, α) with $\alpha = 1$.*

Proof. Necessity is obvious. For sufficiency, let $\mathcal{M}$ be a finite set of weak orderings for a multi-utility representation of $\succsim$ and let μ be the uniform measure on $\mathcal{M}$. Clearly, $(\mu, 1)$ is a confidence model for $\succsim$. ■

In the context of preferences over lotteries, Dubra et al. (2004) show that any reflexive and transitive preference satisfying continuity and independence has an expected multi-utility representation, in which

$$p \succsim q \quad \Leftrightarrow \quad u \cdot p \geq u \cdot q \text{ for all } u \in \mathcal{M} \subseteq \mathcal{U}$$

where $\mathcal{M}$ is closed and convex. The interpretation is the same as that for the finite case, and once again we can relate such preferences to confidence models with a threshold of one. A confidence preference on $\Delta(X)$ represented by a confidence model with a threshold $\alpha = 1$ is transitive. In addition,

Corollary 3. *Any preference $\succsim$ on $\Delta(X)$ satisfying the axioms of Dubra et al. (2004) has a confidence representation (μ, α) where $\mu \in \Delta(\mathcal{U})$ is countably additive and $\alpha = 1$.*

Proof. As before, let $\mathcal{M} \subseteq \mathcal{U}$ be an expected multi-utility representation for $\succsim$, and let m be a strictly-positive, countably-additive measure on $\mathcal{M}$.⁶ If $p \succsim q$, then $u \cdot p \geq u \cdot q$ for all $u \in \mathcal{M}$, and $m(\{u \in \mathcal{M} \mid u \cdot p \geq u \cdot q\}) = 1$. On the other hand, if $\neg(p \succsim q)$, then there exists $u \in \mathcal{M}$ such that $u \cdot p < u \cdot q$, which implies that there exists a relatively open set $U \subseteq \mathcal{M}$ such that $u \cdot p < u \cdot q$ for all $u \in U$. This implies that $m(\{u \in \mathcal{M} \mid u \cdot p \geq u \cdot q\}) < 1$, showing that $(m, 1)$ is a confidence model for $\succsim$. ■

Although the results in Section 3 did not provide a characterization of confidence models on the domain $\mathcal{L} = \Delta(X)$, the preceding corollary establishes that such models do in fact nest expected multi-utility preferences, and clearly they nest traditional EU preferences.

⁶Since $\mathcal{M}$ is closed and hence compact, such a measure exists; consult Corollary 2.8 of Herbert and Lacey (1968).

5.2 Stochastic Preference

There is a long tradition in the literature on stochastic choice⁷ of using a random choice rule ρ to define a stochastic preference $\succsim_\rho$ by

$$x \succsim_\rho y \quad \Leftrightarrow \quad \rho^{\{x,y\}}(x) \geq \rho^{\{x,y\}}(y).$$

To connect this idea with confidence models, let X be a finite set and ρ a RCR on X satisfying the BM-inequalities. Clearly, the preference $\succsim_\rho$ is complete, so by Theorem 1(A) it has a confidence representation. If elements of binary menus are never chosen with probability $1/2$ according to ρ —that is, if $\rho^{\{x,y\}}(x) \neq 1/2$ for all $x, y \in X$ —then $\succsim_\rho$ is antisymmetric and by Theorem 1(B) has a regular confidence representation. In the latter case, it is clear that $\succsim_\rho$ and ρ are compatible on binary sets, so any $\mu \in \Delta(\mathcal{R})$ that represents ρ can be paired with an appropriate $\alpha \in (1/2, 1]$ to form a regular confidence representation for $\succsim_\rho$. However, if there are elements x and y such that $\rho^{\{x,y\}}(x) = 1/2$, then $\succsim_\rho$ is not antisymmetric. Furthermore, $\succsim_\rho$ and ρ will fail to satisfy generalized binary compatibility, and a measure μ representing ρ *cannot* form part of a confidence representation for $\succsim_\rho$. This situation arises due to the requirement for confidence models that $\alpha > 1/2$, which does not prevent confidence models from nesting complete preferences, but does constrain the stochastic choice data that is consistent with such preferences.

5.3 Incomplete Preferences on Anscombe-Aumann Acts

There are certainly domains other than finite sets or lotteries over finite sets on which confidence models might prove interesting. On infinite sets without a lottery structure, one could study confidence models featuring measures on the set of continuous utility functions. But perhaps the most interesting domain to consider is the set of Anscombe-Aumann acts on a finite state space. The classic model for incomplete preferences on this domain is that of Bewley (2002), in which a DM considers a set $\mathcal{M}$ of probability measures on the state space and expresses a preference $f \succsim g$ iff $E_p[u \circ f] \geq E_p[u \circ g]$ for every measure $p \in \mathcal{M}$ and a fixed expected utility function u . Confidence models could generalize this representation by allowing for probabilistic uncertainty over $\mathcal{M}$, and perhaps over u as well.⁸

Confidence models on this domain could potentially be used to address interesting questions on the arrival of information and the updating of incomplete preferences. In addition, confidence models could connect incomplete preferences over acts to random choice over acts, as studied recently by Lu (2013).

⁷See Fishburn (1998) for an overview of this literature.

⁸Ok et al. (2012) and others modify the Bewley model to accommodate a set of utility functions.

Minardi and Savochkin (2013) work with a “graded preference relation” μ where $\mu(f, g)$ is interpreted as the DM’s confidence that f is better than g , and introduce a rule for connecting this primitive to deterministic choice. Comparing and contrasting this approach with the link between incomplete preferences and stochastic choice highlighted in this paper may prove to be an interesting exercise.

A Proofs

A.1 Proof of Corollary 1

We begin by proving Theorem 1(B). To show necessity, let $\succsim$ be represented by (μ, α) for $\mu \in \Delta(\mathcal{R})$ and $\alpha \in (1/2, 1]$. Obviously, $\succsim$ is reflexive: every $R \in \mathcal{R}$ is reflexive. Every $R \in \mathcal{R}$ is also antisymmetric, so if $\mu(\{R \mid xRy\}) \geq \alpha$ then $\mu(\{R \mid yRx\}) \leq 1 - \alpha < 1/2$, showing that $\succsim$ is antisymmetric.

For sufficiency, let $\succsim$ be reflexive and antisymmetric. Let $\varepsilon = [2\binom{N}{N-2}]^{-1}$ and construct a hypothetical RCR ρ on X as follows. On binary sets, define

$$\rho^{\{x,y\}}(x) = \begin{cases} 1/2 + \varepsilon & \text{if } x \succsim y; \\ 1/2 - \varepsilon & \text{if } y \succsim x; \\ 1/2 & \text{otherwise.} \end{cases}$$

For all other sets $D \in \mathcal{D}$, let $\rho^D(x) = 1/|D|$ for all $x \in D$. Recall that the BM inequalities demand

$$BM(\rho, D, x) := \sum_{C: D \subseteq C} (-1)^{|C \setminus D|} \rho^C(x) \geq 0 \quad \forall (D, x) \in \mathcal{D} \times X.$$

The hypothetical ρ constructed above satisfies these inequalities. To see that, note for $|D| \geq 3$, $\rho^D(x) = m(\{R \mid xRy\})$, where m is the uniform measure on $\mathcal{R}$. Thus, letting ρ_m denote the RCR that maximizes the RUF m , we have $BM(\rho, D, x) = BM(\rho_m, D, x) \geq 0$

whenever $|D| \geq 3$. When $|D| = 2$,

$$\begin{aligned}
BM(\rho, D, x) &\geq 1/2 - \varepsilon + \sum_{C: D \subsetneq C} (-1)^{|C \setminus D|} \rho^C(x) \\
&= 1/2 - \varepsilon + \sum_{i=1}^{N-2} \frac{(-1)^i}{i+2} \binom{N-2}{i} \\
&= \left[2 \binom{N}{N-2} \right]^{-1} - \varepsilon \\
&= 0.
\end{aligned}$$

It remains to check $BM(\rho, D, x) \geq 0$ for $|D| = 1$. However, this turns out to be a consequence of the fact that $BM(\rho, D, x) \geq 0$ for $|D| = 2$, as shown by the following lemma from McClellon (2015).

Lemma 3. *Let ρ be an RCR on a set X with $|X| = N$. If $BM(\rho, D, x) \geq 0$ for all pairs (D, x) with $|D| = 2$, then $BM(\rho, \{x\}, x) \geq 0$.*

Applying this result, it follows that the hypothetical ρ constructed using $\succsim$ satisfies the BM-inequalities. Furthermore, it is obvious from the construction of ρ that $\succsim$ and ρ are compatible on binary sets. Therefore, by Lemma 1, $\succsim$ and ρ are consistent, which by Definition 2 entails the existence of a regular confidence model representing $\succsim$.

We now turn to proving Theorem 1(A). Necessity is obvious. Given a reflexive preference $\succsim$, construct the following hypothetical RCR: for $|D| \neq 2$, let $\rho^D(x) = 1/|D|$ for all $x \in D$. For binary sets, once again define $\varepsilon = \left[2 \binom{N}{N-2} \right]^{-1}$ and set

$$\rho^{\{x,y\}}(x) = \begin{cases} 1/2 + \varepsilon & \text{if } x \succsim y \text{ and } \neg(y \succsim x); \\ 1/2 - \varepsilon & \text{if } y \succsim x \text{ and } \neg(x \succsim y); \\ 1/2 & \text{otherwise.} \end{cases}$$

The same arguments used in the proof of Theorem 1(A) show that ρ satisfies the BM-inequalities, and is therefore represented by a RUF $\mu \in \Delta(\mathcal{R})$.

Write $x \sim y$ if $x \succsim y$ and $y \succsim x$. Let I denote the collection of binary sets $\{x, y\}$ with $x \sim y$ and let M denote the number of indifferent pairs in I . For each such pair, define a set $S_{x,y} \subseteq \mathcal{W}$ consisting of all the weak orders W for which (1) xWy and yWx and (2) for all binary sets $\{x', y'\} \neq \{x, y\}$, either $\neg(x'Wy')$ or $\neg(y'Wx')$. In words, $S_{x,y}$ consists of all weak orders for which x and y are the only distinct alternatives that are indifferent to each other.

For each $\{x, y\} \in I$, let $\nu_{x,y}$ denote the uniform measure on $S_{x,y}$. Choose your favorite $\lambda \in (0, 1)$ and define a measure $\mu' \in \Delta(\mathcal{W})$ by

$$\mu' = (1 - \lambda)\mu + \sum_{\{x,y\} \in I} \frac{\lambda}{M} \nu_{x,y}.$$

Observe that if x' and y' are unranked according to $\succsim$, then

$$\mu(\{W \mid x'Wy'\}) = \nu_{x,y}(\{W \mid x'Wy'\}) = 1/2$$

for any $\{x, y\} \in I$, showing that $\mu'(\{W \mid x'Wy'\}) = 1/2$. If $x' \succsim y'$ but $\neg(y' \succsim x')$, then

$$\mu(\{W \mid x'Wy'\}) = 1/2 + \varepsilon \text{ while } \nu_{x,y}(\{W \mid x'Wy'\}) = 1/2$$

again for any $\{x, y\} \in I$. Therefore, $\mu'(\{W \mid x'Wy'\}) = \lambda(1/2 + \varepsilon) + (1 - \lambda)(1/2)$. Similarly, $\mu'(\{W \mid y'Wx'\}) = \lambda(1/2 - \varepsilon) + (1 - \lambda)(1/2)$.

Finally, if $x' \sim y'$ then we have

$$\mu(\{W \mid x'Wy'\}) = \nu_{x,y}(\{W \mid x'Wy'\}) = 1/2$$

provided that $\{x, y\} \neq \{x', y'\}$, and $\nu_{x',y'}(\{W \mid x'Wy'\}) = 1$. Therefore, $\mu'(\{W \mid x'Wy'\}) = (1 - \lambda/M)(1/2) + \lambda/M$. If we set

$$\alpha := \min\{\lambda(1/2 + \varepsilon) + (1 - \lambda)(1/2), (1 - \lambda/M)(1/2) + \lambda/M\} > 1/2$$

then the preceding analysis shows that (μ', α) is a confidence model for $\succsim$.

A.2 Proof of Lemma 2

Let $\succsim$ be a regular confidence preference on $\Delta(X)$. We begin by establishing a simple necessary condition for $\succsim$ to have a regular confidence representation:

Claim: If $\neg(p \succsim q)$, then $\exists \lambda \in (0, 1)$ such that $\neg(\lambda p + (1 - \lambda)q \succsim q)$.

Proof. Let (μ, α) be the regular confidence representation for $\succsim$. Let ρ be the unique RCR on $\Delta(X)$ that maximizes μ . Then we have $\rho^{\{p,q\}}(p) < \alpha$. Let $p\lambda q := \lambda p + (1 - \lambda)q$. By mixture continuity of ρ , there exists $\lambda \in (0, 1)$ such that $\rho^{\{p\lambda q, q\}}(p\lambda q) < \alpha$. The conclusion follows. ■

Let $W := \{(p, q) \in \Delta(X)^2 \mid p \succsim q\}$ and define $\alpha = \inf\{\rho^{\{p,q\}}(p) \mid (p, q) \in W\}$. From this

definition and binary compatibility, we have

$$\rho^{\{p,q\}}(p) \geq \alpha \geq \rho^{\{p',q'\}}(p') \quad (1)$$

for all pairs of lotteries (p, q) and (p', q') with $p \succsim q$ and $\neg(p' \succsim q')$. We claim that in fact the latter inequality holds strictly. If α is attained on W (e.g., in the trivial case of $\alpha = 1$) this follows directly from binary compatibility. In general, additional arguments are needed. So, let $\alpha < 1$ and assume for a contradiction that there exists a pair of lotteries (p', q') with $\neg(p' \succsim q')$ but $\rho^{\{p',q'\}}(p') = \alpha$. By the claim above, there exists some $\lambda \in (0, 1)$ such that $\neg(\lambda p' + (1 - \lambda)q' \succsim q')$. By linearity of ρ , we have

$$\begin{aligned} \rho^{\lambda\{p',q'\}+(1-\lambda)\{q'\}}(\lambda p' + (1 - \lambda)q') &= \lambda \rho^{\{p',q'\}}(p') + (1 - \lambda) \rho^{q'}(q') \\ &= \lambda \cdot \alpha + (1 - \lambda) \cdot 1 =: \beta > \alpha. \end{aligned}$$

By the definition of α , however, there exists some pair $(\hat{p}, \hat{q})$ with $\hat{p} \succsim \hat{q}$ and $\rho^{\{\hat{p},\hat{q}\}}(\hat{p}) < \beta$, which contradicts binary compatibility. Therefore, Equation (1) becomes:

$$\rho^{\{p,q\}}(p) \geq \alpha > \rho^{\{p',q'\}}(p'). \quad (2)$$

Since ρ satisfies GP1–4 there exists a unique ν that maximizes ρ . By definition, this means that for any lotteries p and q we have $\nu(\{u \mid u \cdot p \geq u \cdot q\}) = \rho^{\{p,q\}}(p)$. By Equation (2), $\rho^{\{p,q\}}(p) \geq \alpha$ if and only if $p \succsim q$ and thus (ν, α) is a confidence model for $\succsim$.

A.3 Proof of Theorem 1

Necessity of the axioms is obvious: $\mu(\{u \mid u \cdot p \geq u \cdot p\}) = 1$ so $p \succsim p$ for any p . For Independence, let $p' = \lambda p + (1 - \lambda)r$ and $q' = \lambda q + (1 - \lambda)r$. Clearly $\mu(\{u \mid u \cdot p \geq u \cdot q\}) = \mu(\{u \mid u \cdot p' \geq u \cdot q'\})$, which implies Independence.

The proof of sufficiency rests on the following lemma.

Lemma 4. *Let $p \neq q \in \mathcal{L}$ and fix $\alpha \in (1/2, 1)$. There exists a finitely-additive probability measure $\mu_{pq} \in \Delta(\mathcal{U})$ that satisfies*

$$\mu_{pq}(\{u \in \mathcal{U} \mid u \cdot p' \geq u \cdot q'\}) = \alpha \text{ if } p' = \lambda p + (1 - \lambda)w \text{ and } q' = \lambda q + (1 - \lambda)w$$

for $\lambda \in (0, 1]$ and $w \in \Delta(x)$, and $\mu_{pq}(\{u \in \mathcal{U} \mid u \cdot p' \geq u \cdot q'\}) = 1/2$ for any other $p' \neq q'$.

Proof. Let $\mathcal{A} \subseteq 2^{\mathcal{U}}$ consist of the empty set, $\mathcal{U}$, and in addition all subsets of $\mathcal{U}$ that can be written in the form $A_{p',q'} := \{u \in \mathcal{U} \mid u \cdot p' \geq u \cdot q'\}$ for some $(p', q') \in \mathcal{L}^2$ with $p' \neq q'$.

Define a set function $\nu: \mathcal{A} \rightarrow [0, 1]$ as follows. Let $\nu(\mathcal{U}) = 1$ and $\nu(\emptyset) = 0$. Let $\nu(A_{p',q'}) = \alpha$ if $p' = \lambda p + (1 - \lambda)w$ and $q' = \lambda q + (1 - \lambda)w$ for some $\lambda \in (0, 1]$ and $w \in \Delta(x)$. For any other pair $p' \neq q'$, let $\nu(A_{p',q'}) = 1/2$.

Let $\{A_1, \dots, A_n\}$ and $\{B_1, \dots, B_m\}$ be finite collections of sets in $\mathcal{A}$. No nontrivial sets in $\mathcal{A}$ overlap. The set $A_{p,q}$ must be covered by $\mathcal{U}$ or at least two sets with measure $1/2$. Therefore,

$$\sum_{i=1}^n \mathbf{1}_{A_i} \leq \sum_{j=1}^m \mathbf{1}_{B_j} \quad \Rightarrow \quad \sum_{i=1}^n \nu(A_i) \leq \sum_{j=1}^m \nu(B_j).$$

The Borel σ -algebra on $S_{r,s}$ contains $\mathcal{A}$, so the desired measure μ_{pq} exists by Theorem 3.2.10 in Rao and Rao (1983). $\blacksquare$

With this result in hand, let $\succsim$ be a preference on $\mathcal{L}$ that is reflexive and satisfies Independence. Let $\mathcal{P} \subseteq \mathcal{L}^2$ consist of all pair of lotteries $(p, q) \in \mathcal{L}^2$ with $p \succsim q$ but $p \neq q$. For each such pair, construct a measure μ_{pq} from the lemma above using a common value α' in every construction.

Let $M = |\mathcal{P}|$. Define a measure $\mu \in \Delta(\mathcal{U})$ by

$$\mu = \sum_{(p,q) \in \mathcal{P}} (1/M) \mu_{pq}$$

and let $\alpha = (1/M)\alpha' + ((M - 1)/M)(1/2)$. By Independence, if $(p', q') \notin \mathcal{P}$, then $\mu_{pq}(\{u \mid u \cdot p' \geq u \cdot q'\}) = 1/2$ for all $(p, q) \in \mathcal{P}$. On the other hand, if $(p', q') \in \mathcal{P}$ then $\mu_{p'q'}(\{u \mid u \cdot p' \geq u \cdot q'\}) = \alpha'$, which implies that $\mu(\{u \mid u \cdot p' \geq u \cdot q'\})$ is at least α . It follows that (μ, α) is a confidence representation for $\succsim$.

B Confidence Models and Continuity

Let $\mathcal{L} = \Delta(X)$ and suppose $\succsim$ is represented by a regular confidence model (μ, α) . In the proof of Lemma 2, this is shown to imply that $\succsim$ satisfies a mild continuity property. However, consider the following standard continuity axiom:

Continuity For any $p \in \mathcal{L}$, the sets $\{q \in \mathcal{L} \mid q \succsim p\}$ and $\{q \in \mathcal{L} \mid p \succsim q\}$ are closed.

In Section 3, μ is required only to be finitely additive. If μ fails to be countably additive, $\succsim$ can fail to satisfy continuity. This is somewhat counterintuitive; it occurs even though each u in the support of μ represents an EU preference that certainly does satisfy continuity. I thank an anonymous referee for suggesting the following example: let μ be the finitely-additive measure constructed in Example S1 of the supplement to Gul and Pesendorfer

(2006). Loosely, this measure supports a single utility function u with a uniform tie-breaking rule. If p is indifferent to q under u , then a sequence of lotteries $p_n \rightarrow p$ can be found with each p_n strictly preferred to q . This will violate continuity when a confidence model is constructed by pairing μ with any $\alpha \in (1/2, 1]$.

If μ is required to be countably additive, the RCR ρ represented by μ is continuous in the Hausdorff topology on $\Delta(X)$ and this problem disappears. As stated in the text, I conjecture that reflexivity, antisymmetry, Independence and Continuity are necessary and sufficient for a preference to have a regular confidence representation featuring a countably additive measure. Conceivably, this could be proved by adapting the proof of Theorem 1 and constructing a hypothetical RCR that is compatible on binary sets with the preference, and then applying Lemma 2. The hypothetical RCR would need to be continuous (rather than just mixture continuous) according to the definition in Gul and Pesendorfer (2006). Alternatively or in addition, one could possibly weaken continuity on $\succsim$ to a form of mixture continuity, construct a hypothetical mixture-continuous RCR, and obtain a finitely-additive confidence representation. However, constructing hypothetical RCRs over lotteries is substantially more complex than constructing them over finite domains.

The strategies described in the previous paragraph would generate regular confidence models. In the context of preferences on $\Delta(X)$, dropping regularity is more complicated than merely removing the antisymmetry axiom. Loosely, the reason is that for a confidence preference with representation (μ, α) to have non-trivial indifference, the measure μ must assign positive probability to a set that does not have full dimension in $\mathcal{U}$. This cannot occur for an arbitrary collection of lower-dimensional sets, and the restrictions may vary for the case of finitely- versus countably-additive measures.

References

- Agranov, M. and Ortoleva, P. (2015). Stochastic Choice and Preferences for Randomization. *Working Paper*.
- Ahn, D. S. and Sarver, T. (2013). Preference for Flexibility and Random Choice. *Econometrica*, 81(1):341–361.
- Barberá, S. and Pattanaik, P. K. (1986). Falmagne and the Rationalizability of Stochastic Choices in Terms of Random Orderings. *Econometrica*, 54(3):707–715.
- Bewley, T. F. (2002). Knightian Decision Theory, Part I. *Decisions in Economics and Finance*, 25:79–110.

- Block, H. D. and Marschak, J. (1960). Random Orderings and Stochastic Theories of Responses. In *Contributions to Probability and Statistics*. Stanford University Press.
- Chateauneuf, A. and Faro, J. H. (2009). Ambiguity through confidence functions. *Journal of Mathematical Economics*, 45(9-10):535–558.
- Dubra, J., Maccheroni, F., and Ok, E. (2004). Expected utility theory without the completeness axiom. *Journal of Economic Theory*, 115(1):118–133.
- Evren, O. and Ok, E. (2008). On the Multi-Utility Representation of Preference Relations. *Working Paper*.
- Falmagne, J. C. (1978). A Representation Theorem for Finite Random Scale Systems. *Journal of Mathematical Psychology*, 18:52–72.
- Fiorini, S. (2004). A short proof of a theorem of Falmagne. *Journal of Mathematical Psychology*, 48(1):80–82.
- Fishburn, P. C. (1998). Stochastic Utility. In Barberá, S., Hammond, P. J., and Seidl, C., editors, *Handbook of Utility Theory*, pages 272–311.
- Gul, F. and Pesendorfer, W. (2006). Random Expected Utility. *Econometrica*, 74(1):121–146.
- Herbert, D. J. and Lacey, H. E. (1968). On Supports of Regular Borel Measures. *Pacific Journal of Mathematics*, 27(1):101–118.
- Lu, J. (2013). Random Choice and Private Information. *Working Paper*, pages 1–88.
- McClellon, M. (2015). Unique Random Utility Representations. *Working Paper*.
- McGarvey, D. C. (1953). A Theorem on the Construction of Voting Paradoxes. *Econometrica*, 21(4):608–610.
- Minardi, S. and Savochkin, A. (2013). A Confidence-Based Decision Rule and Ambiguity Attitudes. *Working Paper*, pages 1–22.
- Ok, E., Ortoleva, P., and Riella, G. (2012). Incomplete Preferences Under Uncertainty: Indecisiveness in Beliefs versus Tastes. *Econometrica*, 80(4):1791–1808.
- Rao, K. P. S. B. and Rao, M. B. (1983). *Theory of Charges*. A Study of Finitely Additive Measures. Academic Press.