Theoretical Economics 8 (2013), 503–534 1555-7561/20130503

Contingent preference for flexibility: Eliciting beliefs from behavior

Philipp Sadowski
Department of Economics, Duke University

Following Kreps (1979), I consider a decision maker who is uncertain about her future taste. This uncertainty leaves the decision maker with a preference for flexibility: When choosing among menus that contain alternatives for future choice, she weakly prefers menus with additional alternatives. Standard representations that accommodate this choice pattern cannot distinguish tastes (indexed by a subjective state space) and beliefs (a probability measure over the subjective states) as different concepts. I allow choice between menus to depend on objective states. My axioms provide a representation that uniquely identifies beliefs, provided objective states are sufficiently relevant for choice. I suggest that this result can provide choice theoretic substance to the assumption, commonly made in the (incomplete) contracting literature, that contracting parties who know each others' ranking of contracts also share beliefs about each others' future tastes in the face of unforeseen contingencies.

Keywords. Preference for flexibility, unique beliefs, unforeseen contingencies, incomplete contracts.

JEL classification. D01, D81, D82, D83, D86.

1. Introduction

The expected utility model of von Neumann and Morgenstern (1944, henceforth vNM) explains choice under risk by considering probabilities and taste (a ranking of outcomes) separately. In the context of choice under subjective uncertainty, the corresponding separation of beliefs and tastes is a central concern. For the extreme case where all subjective uncertainty can be captured by objective states of the world, the works of Savage (1954) and Anscombe and Aumann (1963, henceforth AA) achieve this separation. In the opposing extreme, where none of the subjective uncertainty can be captured by objective states, uncertainty can be modeled with a subjective state space. Kreps (1979, henceforth Kreps) and Dekel et al. (2001, henceforth DLR; a relevant corrigendum is Dekel et al. 2007; henceforth DLRS)[1] find that the separation is not possible in this case. This is the standard indeterminacy of state-dependent expected utility models.

Philipp Sadowski: p.sadowski@duke.edu
I thank Roland Benabou and Wolfgang Pesendorfer for their invaluable advice. I am also grateful to David Dillenberger, Faruk Gul, R. Vijay Krishna, Bart Lipman, Mark Machina, Eric Maskin, Todd Sarver, Tymon Tatur, Justin Valasek, numerous seminar audiences, and four anonymous referees for helpful suggestions.

[1] Throughout the paper, I refer to the version of their model that represents preference for flexibility.

DOI: 10.3982/TE884

In the general case, some, but potentially not all, subjective uncertainty can be captured by objective states. This paper analyzes a model of choice under such general subjective uncertainty, which features the AA and DLR models as special cases.[2] The model separately identifies tastes and beliefs over those tastes, provided that objective states are "relevant enough." The main identification result provides a tight behavioral characterization of relevant enough.

The timing of choice is as follows: In period 1, the decision maker (DM) chooses an opportunity act. An opportunity act specifies a menu of alternatives for future choice contingent on the objective state. Between periods 1 and 2, an objective state realizes. In period 2, the act is evaluated and DM gets to choose from the resulting menu. Only period 1 choice is observed. If objective states do not account for all subjective uncertainty that resolves between periods 1 and 2, then DM has contingent uncertainty about her future taste. In that case, commitment to a plan of period 2 choice contingent only on the objective state is costly, and one should observe contingent preference for flexibility: All else being equal, DM prefers an act that assigns a menu that contains more alternatives in any particular state.

This paper provides a representation of such preferences, labeled a representation of contingent preference for flexibility (CPF). Subjective uncertainty that is not captured by the objective state is modeled, as in DLR, via a subjective state space, which collects all possible tastes that might govern DM's choice in period 2. I call it the taste space. The decision maker behaves as if the objective state may be informative about her future tastes, and so conditions her beliefs about future tastes on the objective state. Contingent on the state, choice over menus has a subjective expected utility representation, as in DLR. I show that a central new axiom, Relevant Objective States (Axiom 1), is equivalent to the unique identification of utilities and conditional beliefs in this representation.

To be more specific, let I be the objective state space. An opportunity act, g, assigns a contingent menu of lotteries over prizes, $g(i)$, to every objective state, i in I. The taste space, S, collects all possible vNM rankings of lotteries over prizes. In the case of finite I, choice over acts has a CPF representation if it can be represented by

$$V(g) = \sum_{i \in I} \phi(i) \left[\int_S \Big(\max_{\alpha \in g(i)} U_s(\alpha) \Big) \, d\mu_i(s) \right],$$

where ϕ is a probability measure on I, the realized vNM utility function U_s represents taste s in S, and $\mu_i(s)$ is a probability measure on S. The representation suggests that while the menu of alternatives DM expects to choose from in stage 2 is determined by the objective state i, she anticipates her utility function to be fully determined by her taste s. She also expects to learn s and i prior to choosing an alternative. The measure ϕ is interpreted as DM's beliefs over I and $\mu_i(s)$ is interpreted as the belief that taste s occurs, contingent on i.

[2]In the Savage and Kreps models there are no objective probabilities, while AA, DLR, and the present paper consider a combination of subjective and objective probabilities.

Theorem 1 takes the CPF representation and the distribution ϕ as given.[3] It establishes that conditional beliefs $\mu_i(s)$ are unique and utilities U_s are unique in an appropriate sense if and only if choice between opportunity acts satisfies the Relevant Objective States axiom. The axiom is formulated in terms of DM's ranking of menus contingent on the objective state. This ranking is derived from her choice over acts. Say that two menus are *equivalent for DM* if for every contingent ranking, the union of those menus is indifferent to each of the menus individually. Objective states are relevant if for any two menus that are not equivalent for DM, there is an objective state contingent on which one is strictly preferred over the other.

Theorem 2 states that choice over opportunity acts has a CPF representation if and only if it satisfies the immediate extensions of the state-dependent AA and DLR axioms. These axioms are necessary for a more general representation, where both beliefs and utilities depend on objective states. Given the more general representation, Theorem 2 therefore implies that the assumption that only beliefs condition on objective states does not constrain period 1 choice.

Even though the model does not capture period 2 choice from a menu explicitly, a researcher may want to forecast period 2 choice behavior. The CPF representation describes choice over menus in period 1, as if the DM held beliefs about the tastes that might govern her period 2 choice. Theorem 1 uniquely identifies those beliefs, which are parameters of the representation, from period 1 choice. Therefore, the natural inductive step is to employ the DM's beliefs about future tastes to forecast period 2 choice behavior. There are good arguments against this inductive step. For example, one could instead make period 2 choice part of the domain, leaving less room for erroneous modeling assumptions.[4] However, an essential reason for the use of scientific models is to make predictions about the world based on limited data. Choice theory is well positioned to supply such models for economic applications: Axiomatization translates limited data (here period 1 choice data) to a model, and identification establishes those parameters of the model (here the beliefs) one which one might base inferences.

Being able to forecast behavior can also be important in strategic situations, for example, when one party's valuation of a contract depends on future actions taken by the other party. As an example that illustrates how the CPF representation can apply to the evaluation of contracts, consider a retailer who writes a contract with her supplier today about tomorrow's order. The demand, s, facing the retailer tomorrow is either high (h) or low (l). Today s is unknown to both parties; tomorrow it is the private knowledge of the retailer. (While demand is observable in many situations, unobservable demand levels here simply serve as convenient labels for the different unobservable profit functions the retailer can conceive of.) The only relevant public information that becomes available tomorrow is consumer confidence, i, a general market indicator, which also is either high (H) or low (L). Thus, a contract, g, can only condition on consumer confidence, not on demand. The most efficient contract might give the retailer some choice

[3]The probability measure ϕ may be objective. If ϕ is subjective as suggested above, it must also be elicited from choice. I address this case in Theorem 3.

[4]Ahn and Sarver (2013) provide a model that connects preference for flexibility in period 1 to choice frequencies in period 2.

of supply quantities, q, contingent on consumer confidence. Consider this type of contract. From the retailer's perspective, the contract is an act in the terminology of this paper. Routinely one can write down the following objective function for the retailer's choice between contracts:

$$V(g) = \sum_{i\in\{H,L\}} \phi(i)\Big[\sum_{s\in\{h,l\}} \mu_i(s) \max_{q\in g(i)} (U_s(q))\Big].$$

First, take consumer confidence, $i \in \{H, L\}$, as given. The retailer can then order any quantity in $g(i)$. If she faces demand $s \in \{h, l\}$ tomorrow, she will choose the quantity q that maximizes her profits, $U_s(q)$. Today she does not know tomorrow's demand, but she can assign probabilities conditional on consumer confidence, $\mu_i(s)$. She values the menu $g(i)$ at its expected value, $\sum_{s\in\{h,l\}} \mu_i(s) \max_{q\in g(i)}(U_s(q))$. Second, she takes an expectation over different levels of consumer confidence according to a probability distribution ϕ. This is an example of a CPF representation.[5]

Section 2 investigates conditions under which beliefs are identified in the example above. In the body of the paper, the objective state space, I, is assumed to be finite. Section 3 lays out the model and establishes Theorems 1 and 2. Section 4 contains Theorem 3, which combines the two results and elicits beliefs ϕ on I from choice. Section 5 discusses related literature. Section 6 comments in more detail on possible implications for contracting. Appendix A provides existence and identification results for the case of a general measurable objective state space. Most proofs of the theorems are relegated to Appendix B.

2. Illustration of identification of beliefs

In this section, I consider three cases of a CPF representation: when none of the subjective uncertainty can be captured by objective states (irrelevant objective states), when all of the subjective uncertainty can be captured by objective states (no preference for flexibility), and the general case, where some, but not all, of the subjective uncertainty can be captured by objective states (preference for flexibility and relevant objective states). To illustrate these cases, I use the setup of the above example, except the final outcomes are lotteries, α, over quantities.

- Irrelevant objective states: Suppose that the retailer's beliefs are independent of consumer confidence; that is, $\mu_H(h) = \mu_L(h) = \mu(h)$ and

$$V(g) = \sum_{i\in\{H,L\}} \phi(i)\Big[\sum_{s\in\{h,l\}} \mu(s) \max_{q\in g(i)} (U_s(q))\Big].^{6}$$

  In this case, her induced ranking of menus is independent of consumer confidence. Hence, for the purpose of identifying beliefs μ, it is without loss of

[5] The CPF representation also evaluates more general contracts, where, contingent on consumer confidence, the retailer is given some choice between lotteries over different quantities. For example, the contract might specify an action that has probabilistic consequences.

[6] Ozdenoren (2002) provides a model that generalizes this example, as discussed in Section 5.

generality to consider only contracts with $g(H)=g(L)$. If g is such a noncontingent contract, then

$$V(g)=\sum_{s\in\{h,l\}}\mu(s)\max_{\alpha\in g(H)}(U_s(\alpha)).$$

This is an example of DLR's representation. To see that beliefs are not identified, consider a different probability distribution $\widehat{\mu}(s)$ on $S=\{h,l\}$ and rescaled utilities $\widehat{U}_s(x)=U_s(x)(\mu(s)/\widehat{\mu}(s))$. Then

$$\sum_{s\in\{h,l\}}\mu(s)\Big(\max_{\alpha\in g(H)}U_s(\alpha)\Big)\equiv\sum_{s\in\{h,l\}}\widehat{\mu}(s)\Big(\max_{\alpha\in g(H)}\widehat{U}_s(\alpha)\Big).$$

This is the fundamental indeterminacy in the Kreps and DLR models and variations of those models.

- No preference for flexibility: Suppose that $\mu_H(h)=1$ and $\mu_L(h)=0$. This is the unique (up to relabeling) specification of beliefs where subjective uncertainty is perfectly captured by the objective states and none of the contingent rankings exhibits preference for flexibility. In this case, it is without loss of generality to identify h with H and l with L, and one can confine attention to contracts with lotteries, instead of menus, as outcomes. If $g(i)=\alpha_i$ is such a fully specified contract, then

$$V(g)=\sum_{i\in\{H,L\}}\phi(i)U_i(\alpha_i).$$

This is an example of AA's state-dependent representation.

- Preference for flexibility and relevant objective states: Last, suppose the retailer believes that the probability of high demand is increasing with consumer confidence; that is, $1>\mu_H(h)>\mu_L(h)>0$. Further suppose that there is another representation of the same ranking of contracts based on beliefs over objective states, $\widehat{\phi}$, conditional beliefs, $\widehat{\mu}_i(s)$, and tastes, $\widehat{U}_s$:

$$\widehat{V}(g)=\sum_{i\in\{H,L\}}\widehat{\phi}(i)\Big[\sum_{s\in\{h,l\}}\widehat{\mu}_i(s)\max_{\alpha\in g(i)}(\widehat{U}_s(\alpha))\Big].$$

Choices V and $\widehat{V}$ have to generate the same ranking of contracts.

Consider two quantities (or degenerate lotteries), q_h and q_l, such that the retailer prefers to receive q_h if demand is high and q_l if demand is low, that is, $U_h(q_h)-U_h(q_l)>0$ and $U_l(q_h)-U_l(q_l)<0$. Slightly abusing notation, I denote a lottery that gives q_h with probability α and q_l with probability $1-\alpha$ by α. I denote the menu that contains lotteries α and β by $\{\alpha,\beta\}$. Suppose for some $\beta<\alpha$ and $\delta,\varepsilon\in(0,1-\alpha)$ the retailer is indifferent between the two contracts

$$g = \begin{pmatrix} \{\alpha+\delta, \beta\} & \text{if } i = H \\ \{\alpha, \beta\} & \text{if } i = L \end{pmatrix}$$

$$g' = \begin{pmatrix} \{\alpha, \beta\} & \text{if } i = H \\ \{\alpha+\varepsilon, \beta\} & \text{if } i = L \end{pmatrix}.$$

That $\beta < \alpha$ implies that α is relevant for the value of these contracts only under taste h. Hence, $g \sim g'$ implies that

$$\phi(H)\mu_H(h)\delta(U_h(q_h) - U_h(q_l)) = \phi(L)\mu_L(h)\varepsilon(U_h(q_h) - U_h(q_l)).$$

An analogous equality must hold for the parameters of $\widehat{U}$. Therefore,

$$\frac{\mu_H(h)}{\mu_L(h)} = \frac{\varepsilon\phi(L)}{\delta\phi(H)} \quad \text{and} \quad \frac{\widehat{\mu}_H(h)}{\widehat{\mu}_L(h)} = \frac{\varepsilon\widehat{\phi}(L)}{\delta\widehat{\phi}(H)}.$$

If probabilities of objective states are objective, that is, $\widehat{\phi} = \phi$, then

$$\frac{\mu_H(h)}{\mu_L(h)} = \frac{\widehat{\mu}_H(h)}{\widehat{\mu}_L(h)},$$

and similarly

$$\frac{\mu_H(l)}{\mu_L(l)} = \frac{\widehat{\mu}_H(l)}{\widehat{\mu}_L(l)}.$$

Since μ and $\widehat{\mu}$ are both probability measures, $\mu_H(h)/\mu_L(h) \neq 1$ immediately implies that $\mu \equiv \widehat{\mu}$. Standard arguments, applied to the comparison of contracts that disagree only under state i, imply that the expected utility functions $\widehat{U}_h$ and $\widehat{U}_l$ can only differ from their respective counterparts U_h and U_l by a common linear transformation and the addition of constants. This argument illustrates how identification relies crucially on the fact that beliefs ϕ over objective states are held fixed. More generally, it makes clear why it is necessary to observe the retailer's choice between contracts (opportunity acts), and not just her ranking of menus contingent on each objective state: the willingness to trade off payoffs across objective states (captured by the indifference between contracts g and g') determines the relative weight assigned to taste h under objective state H versus L.[7]

The above reasoning can be generalized to any finite state space, I. If a CPF representation has the feature that there are at least as many linearly independent probability measures over the taste space, indexed by $i \in I$, as there are relevant tastes, then beliefs are uniquely identified and the scaling of utilities is uniquely identified up to a common linear transformation. The next section develops a general model of choice between opportunity acts and characterizes, in terms of behavior, all preferences that have a CPF representation with this feature.

[7]More specifically, my identification strategy relies on the linear aggregation of objective states. Ozdenoren (2002) provides a representation of preferences over opportunity acts that can accommodate ambiguity aversion with respect to objective states.

3. A model with unique beliefs

Consider a two-stage choice problem, where an objective state realizes between the two stages. In period 2, DM chooses a lottery over prizes. In period 1, DM chooses an opportunity act. Such an act is a state contingent specification of a set of lotteries (a menu) that contains the feasible alternatives for period 2 choice.

Let Z be a finite prize space with cardinality k and typical elements x, y, z, and let $\Delta(Z)$ be the space of all lotteries over Z with typical elements α, β, γ. When there is no risk of confusion, x also denotes the degenerate lottery that assigns unit weight to x. Let $\mathcal{A}$ be the collection of all compact subsets of $\Delta(Z)$ with menus A, B, C as typical elements.[8] Further, let I be a finite objective state space with typical elements i, j. Let $\mathcal{F}$ be the σ-algebra generated by the power set of I, where $i, j \in I$ also denote elementary events.[9]

Let G be the set of all opportunity acts with typical elements g, h. An opportunity act is a measurable function $g: I \to \mathcal{A}$. If state i realizes, DM gets to choose an alternative from the menu $g(i) \in \mathcal{A}$. This choice is not modeled explicitly. Instead, $\succ$ is a binary relation on $G \times G$; $\succcurlyeq$ and $\sim$ are defined in the usual way.

The following concepts are important throughout the paper.

Definition 1. The convex combination of menus is defined as

$$pA + (1-p)B := \{p\alpha + (1-p)\beta \mid \alpha \in A, \beta \in B\}.$$

The convex combination of opportunity acts is defined such that

$$(pg + (1-p)h)(i) := pg(i) + (1-p)h(i).$$

To define DM's induced ranking of menus A and B contingent on state $i \in I$, consider acts g_i^A and g_i^B that give menu A and B, respectively, in state i and some arbitrary but fixed default menu, A^*, in every other state. Comparing g_i^A and g_i^B induces a ranking $\succ_i$ over menus. In the context of the model below, $\succ_i$ is independent of A^*.

Definition 2. Fix an arbitrary menu $A^* \in \mathcal{A}$. For $i \in I$ and $A \in \mathcal{A}$, define g_i^A by

$$g_i^A(j) := \begin{cases} A & \text{for } j = i \\ A^* & \text{otherwise.} \end{cases}$$

Let the *contingent ranking* $\succ_i$ be the induced binary relation on $\mathcal{A} \times \mathcal{A}$, $A \succ_i B$ if and only if $g_i^A \succ g_i^B$; $\succcurlyeq_i$ and $\sim_i$ are defined in the usual way. A state $i \in I$ is *nonnull* if there are $A, B \in \mathcal{A}$ with $A \succ_i B$.

In period 2, objects of choice are lotteries over the prize space. The taste space (the collection of all conceivable period 2 tastes) is the collection of all vNM rankings of lotteries. The following definition is due to DLRS.

[8] Compactness is not essential. If menus were not compact, maximum and minimum would have to be replaced by supremum and infimum, respectively, in all that follows.

[9] The case of a general measurable space $(I, \mathcal{F})$ is relegated to Appendix A.

DEFINITION 3. The set

$$S=\left\{s\in\mathbb{R}^k \;\middle|\; \sum_t s_t=0 \text{ and } \sum_t s_t^2=1\right\}$$

is the *taste space*.[10] Let $\mathcal{B}$ be the Borel σ-algebra on S.

The set S collects all possible realized vNM utilities, twice normalized. Every taste in S is a vector with k components, where each entry can be thought of as specifying the relative utility associated with the corresponding prize.[11]

DEFINITION 4. Call (ϕ,μ,U) a *contingent preference for flexibility* (CPF) *representation* of the preference relation $\succ$ if ϕ is a probability measure on I, $\mu=\{\mu_i\}_{i\in I}$ is a collection of probability measures on $(S,\mathcal{B})$, and $U=\{U_s\}_{s\in S}$ is a family of vNM utilities on $\Delta(Z)$, integrable in s, where U_s represents taste s and the objective function

$$V(g)=\sum_{i\in I}\phi(i)\left[\int_S\left(\max_{\alpha\in g(i)} U_s(\alpha)\right)d\mu_i(s)\right]$$

represents $\succ$.

If U_s is a vNM representation of taste s, then it must have the form $U_s(\alpha)=l(s)(s\cdot\alpha)+b_s$, where $s\cdot\alpha$ is the dot product of state s and lottery α, $l(s)$ is the "intensity" of taste s, and b_s is a constant. The relative intensity of utilities together with beliefs determines how DM trades off gains across tastes. The constants b_s have no behavioral content. This motivates the next definition.

DEFINITION 5. Let (ϕ,μ,U) be a CPF representation of $\succ$.

(i) The *space of relevant objective states*, $I^*\subseteq I$, is the minimal set with $\phi(I^*)=1$. The *space of relevant tastes* is $S^*:=\bigcup_{i\in I^*}\operatorname{supp}(\mu_i)$.[12]

(ii) Beliefs μ and tastes U are *unique given* ϕ, if for any other CPF representation $(\phi,\widehat{\mu},\widehat{U})$ of $\succ$, there are $a>0$ and $\{b_s\}_{s\in S^*}\subset\mathbb{R}$ such that for all $i\in I^*$ and all $S'\subseteq S^*$, $\widehat{\mu}_i(S')=\mu_i(S')$ and $\widehat{U}_s=aU_s+b_s$ holds μ_i-almost everywhere.

The set S^* can be thought of as the set of tastes DM considers possible. An axiomatization of the CPF representation is given in Theorem 2. The distribution ϕ is identified from behavior in Theorem 3. The main concern, however, is to separately identify beliefs μ and tastes U, *provided* that DM's choice over acts has a CPF representation for a given distribution ϕ.

[10] DLRS refer to S as the universal state space.

[11] In the context of the representation theorem in DLRS, as in the theorems that follow, there is clearly always a larger taste space, also allowing a representation of $\succ_i$, in which multiple tastes represent the same ranking of lotteries.

[12] The support of a measure is the closure of the collection of points for which every neighborhood in $\mathcal{B}$ has positive measure.

The main new axiom of this paper can be paraphrased as follows: if two menus are not *equivalent for DM*, in the sense that they provide her with the same utility under every relevant taste, then there exists an objective state contingent on which one is preferred over the other.

Axiom 1 (Relevant objective states). *If $A \cup B \nsim_i B$ for some $i \in I$, then there is $j \in I$ with $A \nsim_j B$.*

Two menus that are distinct elements of $\mathcal{A}$ might still be equivalent for DM. If A and B are equivalent for DM, then she should be willing to choose from $A \cup B$ by simply ignoring A.[13] This cannot be the case if $A \cup B \nsim_i B$ for some $i \in I$. If also $A \nsim_i B$, then Axiom 1 is empty. If $A \sim_i B$, then $A \cup B \nsim_i B$ implies that, contingent on i, the item chosen from $A \cup B$ must sometimes be in A and sometimes in B. Axiom 1 requires that there exists a contingent ranking for which either one or the other case becomes more important, namely that there is $j \in I$ with $A \nsim_j B$.[14]

Axiom 1 is not a strong assumption in the sense that it is local; it requires only breaking indifference. For comparison, AA require that there is no relevant subjective uncertainty, contingent on the state of the world. That is, choice between menus would have to satisfy state contingent strategic rationality: If $A \cup B \succ_i B$, then $A \sim_i A \cup B$.[15] In terms of the example from the Introduction, consumer confidence (the objective state) may be relevant for the retailer's beliefs about her desire to order a large or a small quantity, but it is conceivable that the retailer prefers a large quantity even when confidence is low and vice versa. This notion is weaker than the assumption of state contingent strategic rationality, according to which the retailer always prefers the large quantity when confidence is high and the small quantity when confidence is low.

Theorem 1. *If $\succ$ has a CPF representation (ϕ, μ, U), then the following statements are equivalent.*

(i) The binary relation $\succ$ satisfies Axiom 1.

(ii) Beliefs μ and tastes U are unique given ϕ.

(iii) The cardinality of S^* equals the number of linearly independent elements in $\{\mu_i\}_{i \in I^*}$.

If a decision maker behaves as if she has preference for flexibility because of uncertainty about her future taste, updates her beliefs over tastes when learning the objective state, and maximizes her expected utility according to objective probabilities over those states, then her preferences satisfy Axiom 1 if and only if her beliefs over future tastes are

[13] Implicit in the interpretation is that, ultimately, only the chosen item matters for the value of a menu.

[14] To see the content of the axiom in terms of the representation, suppose for simplicity that there are only two subjective states, s_A and s_B, where in state s_A the best element of $A \cup B$ lies in A, and in state s_B it lies in B. Suppose both states are supported by μ_i. This immediately implies $A \cup B \nsim_i B$. If $A \sim_i B$, then the axiom requires that there is $j \in I$ with $\mu_j(s_A) \neq \mu_i(s_A)$.

[15] Axiom 1 is immediately satisfied: $A \cup B \nsim_i B$ implies $A \nsim_i B$.

determined uniquely. This identification gives meaning to the description of beliefs and tastes as distinct concepts. Lack of this distinction is the central drawback of previous work on preference for flexibility, starting with Kreps.

Another difficulty in the application and interpretation of models of preference for flexibility is the infinite subjective state space. Theorem 1(iii) conveniently constrains the space of relevant tastes, S^*, to be smaller than I, which is finite. Axiom 1 implies this finiteness, because I must be rich enough to distinguish between any two menus for which DM might have preference for flexibility. This implies that only finitely many lotteries can be appreciated in any menu. Theorem 1′ in the Appendix A generalizes the result and considers I to be a general topological space, lifting the constraint on the cardinality of S^*.

Finally, given a particular CPF representation of $\succ$, Theorem 1(iii) provides a criterion to check whether $\succ$ satisfies Axiom 1. This criterion is illustrated in the example in Section 2.

To see how relevant objective states imply unique beliefs and utilities, fix the distribution of objective states, ϕ, and suppose there are two CPF representations of the same preference relation, (ϕ, μ, U) and $(\phi, \widehat{\mu}, \widehat{U})$ with corresponding value functions $V(g) = \sum_{i\in I} \phi(i)V_i(g(i))$ and $\widehat{V}(g) = \sum_{i\in I} \phi(i)\widehat{V}_i(g(i))$, where $V_i(A) = \int_S \max_{\alpha\in A}(U_s(\alpha))\, d\mu_i(s)$ and analogously for $\widehat{V}_i(A)$. Neglecting additive constants, additive separability of the representations implies that $\phi(i)V_i(\cdot) = \lambda\phi(i)\widehat{V}_i(\cdot)$ for all $i \in I^*$ and for some $\lambda > 0$. Suppose further that for the contingent ranking $\succ_i$, one could construct menus $K \sim_i \widehat{K}$ such that K generates constant utility payoffs across tastes according to (ϕ, μ, U) and $\widehat{K}$ does so according to $(\phi, \widehat{\mu}, \widehat{U})$. On the one hand, $K \sim_i \widehat{K}$ implies $V_i(K) = V_i(\widehat{K}) = \lambda\widehat{V}_i(\widehat{K})$. On the other hand, changing the objective state from i to j changes only DM's beliefs about her future tastes. If a menu generates the same utility payoff for every taste, then the conditional value of the menu is independent of the objective state: $V_j(K) = V_i(K)$ and $\widehat{V}_j(\widehat{K}) = \widehat{V}_i(\widehat{K})$ for all $j \in I^*$. Hence, $V_j(K) = \lambda\widehat{V}_j(\widehat{K})$ or $K \sim_j \widehat{K}$ would have to hold for all $j \in I^*$. At the same time, if (ϕ, μ, U) and $(\phi, \widehat{\mu}, \widehat{U})$ were distinct, $\widehat{K}$ would not generate constant utility payoffs across tastes according to (ϕ, μ, U), because utility payoffs depend on the intensities of U and $\widehat{U}$, respectively. Therefore, $K \cup \widehat{K} \succ_{j'} K$ for some $j' \in I$. Axiom Relevant Objective States then implies that there is $j \in I$ with $K \nsim_j \widehat{K}$, a contradiction. This rough intuition does not quite work, because the construction of menus that generate the same utility payoff for every taste is not always possible. Because $S^* \subset S$ is finite, however, one can construct pairs of menus (A, B for (ϕ, μ, U) and $\widehat{A}, \widehat{B}$ for $(\phi, \widehat{\mu}, \widehat{U})$) for which the difference in utility payoffs is constant across tastes. Let K be the convex combination of menus $(1/2)A + (1/2)\widehat{B}$ and let $\widehat{K} = (1/2)\widehat{A} + (1/2)B$. Then $K \sim_i \widehat{K}$ implies that $K \sim_j \widehat{K}$ for all $j \in I$. By the type of argument laid out above, $K \cup \widehat{K} \succ_{j'} K$ for some $j' \in I$. This contradicts Axiom 1.

Ozdenoren (2002) analyzes the case where Axiom 1 fails completely, in the sense that objective states are irrelevant to the decision maker. Then only the support of the probability measures μ_i that allow a representation can be identified. This is the same indeterminacy encountered in DLR.

In addition to the key role that objective states play in Axiom 1, the fact that menus consist of lotteries is also important for the identification of beliefs. Nehring (1999) finds that acts with menus of prizes as outcomes do not allow the separate identification of tastes and beliefs in the axiomatic setup developed by Savage (1954). To establish the uniqueness result, the payoff generated by a menu must be varied independently for different tastes. This is possible only because DM can be offered lotteries over prizes.

Remark. A remark on the interpretation of tastes, or subjective states, is in order. Suppose for a moment that there is an underlying state space Ω, that provides a complete description of all relevant aspects of the world. That is, $\omega \in \Omega$ even determines DM's taste, $s \in S$. In that case, S generates a sub-σ-algebra on Ω. The question is the extent to which Ω is *observable*. Let I be the collection of observable events $i \subset \Omega$, where I generates another sub-σ-algebra on Ω. Now consider a probability measure μ on Ω that represents DM's beliefs. If there is no correlation between events in I and events in S, then the induced marginal distribution $\mu_i(s)$ is independent of i, and the objective state space Ω can be dropped from the description of the model, as in DLR. For example, Ω could be the product space $I \times S$ and μ could be a product measure. If, to the other extreme, there is perfect correlation between events in I and events in S, then I itself can play the role of the complete objective state space in (the state-dependent version of) the AA model. Theorem 1 is concerned with the general case of some, but not perfect, correlation. The examples in Section 2 illustrate the three cases.

I now establish existence of a CPF representation. As mentioned above, the axioms are direct extensions of familiar assumptions.

Axiom 2 (Preference). *The binary relation $\succ$ is asymmetric and negatively transitive.*

Axiom 3 (vNM Continuity). *If $g \succ h \succ g'$, then there exist $p, q \in (0,1)$ such that $pg + (1-p)g' \succ h \succ qg + (1-q)g'$.*

Axiom 4 (Independence). *If $g \succ g'$ for $g, g' \in G$ and if $p \in (0,1)$, then*

$$pg + (1-p)h \succ pg' + (1-p)h$$

for all $h \in G$.

If a convex combination of menus were defined as a lottery over menus, then the motivation of Independence in my setup would be the same as in more familiar contexts. Uncertainty would resolve before DM consumes an item from one of the menus. However, following DLR and Gul and Pesendorfer (2001), I define the convex combination of menus as the menu that contains all the convex combinations of their elements. The uncertainty generated by the convex combination is only resolved after DM chooses an item from this new menu. Gul and Pesendorfer term the additional assumption needed to motivate Independence in this setup "indifference as to when uncertainty is resolved."

Axiom 5 (Nontriviality). *There are $g, h \in G$ such that $g \succ h$.*

The next axiom considers DM's contingent ranking of menus, $\succ_i$. As long as some subjective uncertainty is not captured by objective states, $\succ_i$ should exhibit preference for flexibility. This is captured by the central axiom in Kreps, which states that larger menus are weakly better than smaller menus.

Axiom 6 (Monotonicity). *For all $A, B \in \mathcal{A}$ and all $i \in I$, $A \cup B \succsim_i A$.*

Lemma 1. *If $\succ$ satisfies Axioms 2–6, then $\succ_i$ is a preference relation and satisfies the appropriate variants of vNM Continuity, Independence, and Monotonicity for all $i \in I$. Furthermore, there is a nonnull event $i \in I$.*

The proof is immediate.

Theorem DLRS (Theorem 2 in DLRS). *For $i \in I$ nonnull, $\succ_i$ is a preference that satisfies the appropriate variants of vNM Continuity, Independence, and Monotonicity if and only if there is a subjective state space S_i, a positive countably additive measure μ_i on S_i, and a set of nonconstant and continuous expected utility functions $U_{s,i} : \Delta(Z) \to \mathbb{R}$ such that*

$$V_i(A) = \int_{S_i} \max_{\alpha \in A} U_{s,i}(\alpha)\, d\mu_i(s)$$

represents $\succ_i$ and every vNM ranking of lotteries in $\Delta(Z)$ corresponds to at most one state in S_i.[16]

Because $U_{s,i}(\alpha)$ are realized vNM utility functions, the subjective state space S_i can be replaced by the taste space S for all $i \in I$. Note that the taste space does not include the taste where DM is indifferent between all prizes, implicitly assuming nontriviality of the ex post preferences over prizes.

Theorem 2. *The binary relation $\succ$ satisfies Axioms 2–6 if and only if it has a CPF representation.*

The proof first employs the mixture space theorem (Kreps 1988, Theorem 5.11) to establish an additively separable representation of $\succ$. That is, $\sum_{i \in I} v_i(g(i))$ represents $\succ$ for some family of utility functions, $\{v_i\}_{i \in I}$, on $\mathcal{A}$, where v_i are unique up to a common positive linear transformation and the addition of constants.[17] Now consider some linear representation, $\widehat{v}_i$, of $\succ_i$ on $\mathcal{A}$. Since v_i also represents $\succ_i$, the mixture space theorem implies that v_i agrees with $\widehat{v}_i$ up to scaling. The scaling is absorbed by $\phi(i)$, which is then normalized to be a probability distribution. Thus,

$$V(g) = \sum_{i \in I} \phi(i) \widehat{v}_i(g(i))$$

[16] See footnotes 3 and 5 in DLRS.

[17] The result is first established for convex valued acts, defined in the Appendix A, and then extended to all acts.

represents $\succ$. Note that this is AA's state-dependent representation, with the exception that opportunity acts have menus as outcomes, while AA acts have lotteries as outcomes. Indeed, Axioms 2–4 imply AA's axioms. Furthermore, Axioms 2–6 imply DLRS' axioms, as shown in Lemma 1. According to Theorem DLRS, $\succ_i$ can then be represented by

$$\widehat{V}_i(A) = \int_S \max_{\alpha \in A}(\widehat{U}_{s,i}(\alpha))\, d\widehat{\mu}_i(s),$$

where $\widehat{\mu}_i$ is a probability measure on S and $\widehat{U}_{s,i}$ is a vNM utility function that represents taste $s \in S$, that is, $\widehat{U}_{s,i}$ and $\widehat{U}_{s,j}$ are identical up to a positive affine transformation. Pick any $j \in I$ and define $U_s := \widehat{U}_{s,j}$. The lack of identification in DLRS implies that there is a measure μ_i on S, such that $\mu_i(s)U_s \propto \widehat{\mu}_i(s)\widehat{U}_{s,i}$. Therefore, $\succ_i$ can be represented by

$$V_i(A) = \int_S \max_{\alpha \in A} U_s(\alpha)\, d\mu_i(s)$$

for all $i \in I$. Since V_i is linear, there is a CPF representation (ϕ, μ, U); that is,

$$V(g) = \sum_{i \in I} \phi(i)\left[\int_S \max_{\alpha \in g(i)} U_s(\alpha)\, d\mu_i(s)\right]$$

represents $\succ$. The intensity of each taste is endogenous, but it is fixed across objective states.

Clearly Axioms 2–6 are also necessary for the generic combination of the AA and the DLRS representations,

$$\widehat{V}(g) = \sum_{i \in I} \phi(i)\widehat{V}_i(g(i)) = \sum_{i \in I} \phi(i)\left[\int_S \max_{\alpha \in g(i)} (U_{s,i}(\alpha))\, d\mu_i(s)\right],$$

where objective states impact not only probabilities, μ_i, but also the intensities of tastes. Theorem 2 implies that there is a CPF representation of $\succ$ whenever the more general representation $\widehat{V}$ exists. Therefore, the assumption that only beliefs condition on objective states does not constrain period 1 choice.

4. Probabilities over objective states

Theorem 1 takes the distribution ϕ on I and a CPF representation (ϕ, μ, U) as given, and establishes that μ and U are unique in the appropriate sense if and only if objective states are relevant. The distribution ϕ might be objective in the sense that it corresponds to observed frequencies of objective states, or it might be subjective, in which case it must also be elicited from behavior. The unique identification of ϕ is analogous to the classical problem addressed by AA. There, the unique identification of probabilities of observable states is based on the assumption of *state independence* of the ranking of outcomes. The difference is that they consider acts with lotteries (instead of menus of lotteries) as outcomes, so there is no room for preference for flexibility in their setup. In my setup, the combination of *objective state independence* and Axiom 1 rules out any preference for flexibility. Thus, the state independence assumption has to be confined to a proper subset $\Psi \subset \mathcal{A}$ to be useful here. Having assumed state-independent rankings, AA consider only cardinally state-independent rankings (or state-independent utilities). This cannot be assumed in terms of an axiom. Instead it is a constraint on the class of representations for which they establish their uniqueness result. For the CPF representation, it would amount to requiring that $\int_S \max_{\alpha \in A} U_s(\alpha)\, d\mu_i(s)$ is independent of $i \in I^*$ for all $A \in \Psi$. But if $\Psi \subset \mathcal{A}$ is a generic collection of menus, then this might not be consistent with $\succ$, which applies to all of G.[18] Thus, the requirement must be confined to a *particular* collection of menus.

Definition 6. Let $X \subseteq Z$ denote a set of prizes and let $\Delta(X)$ denote the set of all lotteries with support in X. Let $\Psi(\Delta(X)) \subseteq \mathcal{A}$ be the set of all menus of lotteries that have support in X.

Definition 7. A CPF representation, (ϕ, μ, U), is *state-independent with respect to* $X \subseteq Z$ if $\int_S \max_{\alpha \in A} U_s(\alpha)\, d\mu_i(s) = \int_S \max_{\alpha \in A} U_s(\alpha)\, d\mu_j(s)$ for all $A \in \Psi(\Delta(X))$ and all $i, j \in I^*$. Further, (ϕ, μ, U) is the *unique CPF representation that is state-independent with respect to X* if for any other CPF representation $(\widehat{\phi}, \widehat{\mu}, \widehat{U})$ that is state-independent with respect to X, $\widehat{\phi} \equiv \phi$, $\widehat{\mu}_i \equiv \mu_i$ for all $i \in I^*$, and there are $a > 0$ and $\{b_s\}_{s \in S^*} \subset \mathbb{R}$ such that $\widehat{U}_s \equiv aU_s + b_s$ for all $s \in S^*$.

Axiom 7 (Partial objective state independence). *There is a nondegenerate $X \subseteq Z$ such that for $A, B \in \Psi(\Delta(X))$, $A \succ_i B$ for some $i \in I$ implies $A \succ_j B$ for all nonnull $j \in I$.*

To illustrate Axiom 7, consider $X = \{\$1, \$0\}$ to consist of the prizes “1 dollar” and “nothing.” The first part of Axiom 7 then requires that the ranking of menus that consist only of lotteries that pay out either \$1 or nothing must be state-independent. To motivate the requirement, it is sufficient to assume that the value of \$1 (versus nothing) is state-independent. The assumption does not require that DM is certain about his taste, but only that the objective state is uninformative about this uncertainty.[19]

Theorem 3. *The binary relation $\succ$ satisfies Axioms 1–7 if and only if there exists some nondegenerate $X \subseteq Z$ such that $\succ$ has a unique CPF representation that is state-independent with respect to X, (ϕ, μ, U). In this representation, $U_s(x)$ is constant across S^* for all $x \in X$.*[20]

Proof. For a CPF representation where $\int_S \max_{\alpha \in A} U_s(\alpha)\, d\mu_i(s)$ does not depend on $i \in I$ for any $A \in \Psi(\Delta(X))$, the uniqueness of ϕ follows in complete analogy to the corresponding result in AA. Given this unique ϕ, Theorem 1 implies uniqueness of μ and

[18] For a simple example of such inconsistency, consider $\Psi = \{\{\alpha\}, \{\beta\}, \{\gamma\}\}$, but for some $p \in (0, 1)$ and $i, j \in I$, $\{p\alpha + (1-p)\gamma\} \succ_i \{\beta\} \succ_j \{p\alpha + (1-p)\gamma\}$. Since $\int_S \max_{\alpha \in A} U_s(\alpha)\, d\mu_i(s)$ is linear, it cannot be independent of $i \in I$.

[19] Naturally, certainty about (some aspects of) tastes also identifies beliefs, as DLR observe. For a brief discussion, see Section 5.

[20] In the case where ϕ is objective, it is possible to strengthen Axiom 7 such that the *unique* CPF representation in Theorem 3 is based on ϕ.

uniqueness of U_s up to a common rescaling and the addition of constants. The existence of a representation where $U_s(x)$ is constant across S for all $x \in X$ is established in Appendix B. Hence, the unique representation must have this feature. □

Since $U_s(x)$ is constant across S^* for all $x \in X$, it follows immediately that there is no preference for flexibility with respect to alternatives in $\Delta(X)$. To see why this must be the case, consider menus $A, B \in \Psi(\Delta(X))$ with $A \sim_i B$ for some $i \in I$. By Axiom 7, $A \sim_j B$ for all $j \in I$. Now suppose that there is preference for flexibility with respect to those menus, $A \cup B \succ_{i'} A$ for some $i' \in I$. By Axiom 1, $A \nsim_j B$ for some $j \in I$, which is a contradiction.[21]

5. Related literature

Ozdenoren (2002) also considers preference for flexibility in the presence of objective states of the world. Instead of Relevant Objective States (Axiom 1), which ensures that contingent rankings are sufficiently different, he assumes that all contingent rankings are the same. Consequently, beliefs over tastes are not identified in his model.

I know of three other identification results that deliver unique beliefs over future tastes for consumption in models of preference for flexibility. First, note that AA's identification of unique beliefs over objective states does not require *full* state independence of preferences.[22] In analogy to AA's argument, beliefs over tastes in the DLR model can be identified uniquely, as long as DM has no preference for flexibility with respect to *part* of the prize space. As an example, DLR suggest to consider a DM without preference for flexibility on one dimension of a product prize space (Shenone 2010 provides details.) Second, Ahn and Sarver (2013) provide a model that requires both choice between as well as random choice from menus to be observable. Their model restricts the beliefs that feature in the representation of choice between menus to correspond to the choice frequencies that describe choice from menus. Finally, in a dynamic model of preference for flexibility, Krishna and Sadowski (2013) show that the DM's attitude toward intertemporal trade-off can also uniquely identify beliefs. They proceed to characterize a behavioral comparison of "greater preference for flexibility" in terms of a stochastic

---

Axiom 7* (Objective probabilities). *There is $X \subseteq Z$ such that for $A, B \in \Psi(\Delta(X))$, and nontrivial D and $D' \in \mathcal{F}$,*

$$\frac{\phi(D')}{\phi(D)+\phi(D')} h_D^A + \frac{\phi(D)}{\phi(D)+\phi(D')} h_{D'}^B \sim \frac{\phi(D)}{\phi(D)+\phi(D')} h_{D'}^A + \frac{\phi(D')}{\phi(D)+\phi(D')} h_D^B.$$

This implies Axiom 8. It also implies that $V(g_D^A) - V(g_D^B) = (V(g_{D'}^A) - V(g_{D'}^B))(\phi(D)/\phi(D'))$ for $A, B \in \Psi(\Delta(X))$.

[21] Once AA restrict attention to representations with state-independent utilities, there is no arbitrariness in their model. In contrast, strict preference for flexibility implies that X is a proper subset of Z. Hence, $\succ$ could satisfy Axiom 7 for some X and Y with $X \neq Y$, but not for $X \cup Y$. Either those menus with support in X or those with support in Y could then be assigned a cardinal ranking, which is state-independent. The two assumptions clearly lead to different representations. The following assumption would rule out this scenario: *If $\succ$ satisfies Axiom 7 for $X \subseteq Z$ and for $Y \subseteq Z$, then it also satisfies the condition for $X \cup Y$.*

[22] This insight also underlies the elicitation of beliefs, ϕ, over objective states in Section 4 of this paper.

dominance condition on the beliefs. Without identification of beliefs, such a comparison is not possible.[23]

The domain of opportunity acts is first analyzed by Nehring (1999), and the notion of contingent menus appears in Epstein (2006). Following Nehring (1996), a companion paper to Nehring (1999), Epstein and Seo (2009) consider a domain of random menus, which are lotteries with menus as outcomes. On this domain they establish unique induced probability distributions over ex post upper contour sets as the strongest possible uniqueness statement.

Theorem 1 does not only provide unique beliefs, but also establishes the finiteness of the collection of relevant tastes, S^*. Dekel et al. (2009) and Kopylov (2009) generate finiteness of S^* in the absence of objective states by basically assuming that the number of lotteries DM can appreciate in any given menu is limited.

Finally, note that the state-independent version of AA's representation can be viewed as a special case of a unique CPF representation, where there is only one taste and the intensity of this one taste is independent of the objective state. Karni and coauthors, for example, Grant and Karni (2005), Karni (2008, 2011a, 2011b), elaborate the point that interpreting AA's or Savage's (1954) unique subjective probabilities of observable states as DM's true beliefs may be misleading, in case the true intensity of her only taste is actually not state-independent. The CPF model is not immune to this concern: Even if choice has a CPF representation, DM's true intensities of tastes might not actually be state-independent. Similarly, the DM might not actually use the expected utility criterion to evaluate uncertain prospects, or alternatives other than the one that is ultimately chosen might also generate utility. None of those modeling assumptions remains innocuous, once beliefs are used to forecast period 2 choice.

6. Asymmetric information and contracts

The CPF model interprets choice between acts in terms of unique state contingent beliefs over future tastes. As I discuss in the Introduction, considering those beliefs as predictors of future choice is an additional assumption.[24] Suppose in a particular situation that it is commonly assumed that the DM's beliefs over tastes provide an accurate forecast of future choice. Conceptually, an observer could agree on those beliefs, as the DM's future choice is potentially observable. The identification of beliefs from behavior is necessary to render such agreement behaviorally meaningful. In this sense, my model makes it possible to talk about "common priors over tastes," just as the identification of beliefs in the work of Savage (1954) makes it possible to talk about common priors over objective states.

The (incomplete) contracting literature routinely assumes that both players have common priors over each other's future tastes.[25] This allows each player to rank all

[23]Limited by the lack of identification in their model, DLR suggest an alternative notion that can be characterized in terms of the support of the beliefs.

[24]That is, beliefs correspond to choice frequencies. The assumption that beliefs are meaningful beyond their role in the representation of individual choice also underlies the notion of "objective probabilities" on which all agents can agree, even if they behave differently.

[25]Section 3 in Maskin and Tirole (1999) elaborates this point.

contracts even if his or her valuation of a contract depends on the other player's future choice. It follows immediately that both players also know each other's ranking of contracts. However, the assumption of commonly known rankings neither implies common beliefs over future tastes nor implies that players know each other's beliefs.

The weaker assumption of commonly known rankings is usually justified by some informal story of learning from past observations. This assumption is not my focus, and I adopt it without doing justice to the game theoretic complexity of the contracting problem. Instead, I address the stronger assumption of common beliefs. This assumption is particularly troubling in the context of indescribable or unforeseen contingencies,[26] where it seems natural that each party has an informational advantage with regard to their own future taste. In a survey on incomplete contracts, Tirole (1999) speculates that

> "...there may be interesting interaction between 'unforeseen contingencies' and asymmetric information. There is a serious issue as to how parties [...] end up having common beliefs ex ante."

My domain is well suited to describe the type of (incomplete) contracts illustrated in the Introduction, where player 1 is given some control rights contingent on observable states, I. For those contracts, the CPF representation gives choice theoretic substance to the assumption of common beliefs[27] and even suggests a possible mechanism for their convergence: If player 1 (the controlling party) is more knowledgeable about her own future taste and if her beliefs can be deduced from her commonly known ranking of contracts, then player 2 should view those beliefs as the true probability distribution of player 1's future taste, and adopt them as his own beliefs when evaluating contracts.

For instance, in the example in the Introduction, suppose the supplier prefers to supply the large quantity to the retailer. He then assigns a higher value to a contract that gives control over the supplied quantity to the retailer if he expects the retailer to order the large quantity more frequently. Further, suppose the retailer is better informed about her own future profit function (or taste) and, hence, the probability with which she will order the large quantity. The supplier would want to learn this information from the retailer before agreeing to a contracts. In contrast, the supplier does not care about the intensity of the retailer's taste. In my axiomatic setup, these two are distinct concepts, and the supplier can elicit the probability distribution over the retailer's future tastes from her ranking of contracts.

Appendix A: Measurable objective state space

If the objective state space I is finite as in the body of the paper, then Axiom 1 limits the cardinality of the space of relevant tastes, S^*. In many standard models, the state

[26] Kreps (1992) points out that a subjective taste space naturally accounts for contingencies that are not just unobservable or indescribable, but unforeseen, at least by the observer.

[27] Dekel et al. (1998) note that

> "...there are very significant problems to be solved before we can generate interesting conclusions for contracting [...] while the Kreps model (and its modifications) seems appropriate for unforeseen contingencies, [...] there are no meaningful subjective probabilities. A refinement of the model that pins down probabilities would be useful."

space is infinite; for example, the objective state space in Savage (1954) and possibly also the space of relevant subjective states in DLR. In most applications, this is a disadvantage, as a finite subjective state space is interpretationally appealing and analytically convenient. However, for those applications where an infinite subjective state space is necessary, my results can be extended to the case when I is infinite, thereby removing the constraint on the cardinality of S^*.

Let $\mathcal{F}$ be a σ-algebra on I. In this context, let G be the collection of simple opportunity acts with typical elements g, h. A simple opportunity act is a measurable function $g: I \to \mathcal{A}$ such that there exists a finite and measurable partition $\{D_t \mid t \in \{1, \dots, T\}\}$ of I with $g(i) = g(j)$ if there is $D \in \{D_t \mid t \in \{1, \dots, T\}\}$ with $i, j \in D$. The operator $\succ$ is a binary relation on $G \times G$.[28] The definition of $\succ_D$ for $D \in \mathcal{F}$ is analogous to the definition of $\succ_i$, Definition 2. Definition 4 of the CPF representation remains valid, where the value function now takes the form

$$V(g) = \int_I \left[\int_S \left(\max_{\alpha \in g(D_t)} U_s(\alpha) \right) d\mu_i(s) \right] d\phi(i),$$

where ϕ is a countably additive probability measure on $(I, \mathcal{F})$, and where μ is a well defined stochastic kernel between $(I, \mathcal{F})$ and $(S, \mathcal{B})$. Definitions and results that generalize those in Section 3 are distinguished by a prime on their label.

Definition 5′. Let (ϕ, μ, U) be a CPF representation of $\succ$.

(i) Let $I^* := \text{supp}(\phi)$ and let $\mathcal{F}^*$ be the σ-algebra on I^* that corresponds to $\mathcal{F}$. The *space of relevant tastes* is $S^* := \bigcup_{D \in \mathcal{F}^*} \text{supp}(\mu_D)$ with Borel σ-algebra $\mathcal{B}^*$.

(ii) Beliefs μ and tastes U are *unique given* ϕ if for any other CPF representation $(\phi, \widehat{\mu}, \widehat{U})$ of $\succ$, the functions $\widehat{\mu}$ and μ induce the same kernel between the measurable spaces $(I^*, \mathcal{F}^*)$ and $(S, \mathcal{B})$, and there is $a > 0$ and an integrable function $b: S^* \to \mathbb{R}$ such that $\int_{S'} \widehat{U}_s \, d\mu_D(s) = \int_{S'} (aU_s + b(s)) \, d\mu_D(s)$ for all $D \in \mathcal{F}^*$ and all $S' \in \mathcal{B}^*$.

The next definition provides a measure of how much set A is preferred over set B in terms of how much the menu corresponding to the entire prize space, Z, is preferred over the worst prize.

Definition 9. Given $D \in \mathcal{F}$, let $\underline{z}$ be the worst prize: $A \succsim_D \{\underline{z}\}$ for all $A \in \mathcal{A}$. For $A, B \in A$, define $p_{A,B}(D) \in (-1, 1)$ such that the following conditions hold.

(i) For $A \succsim_D B$, $p = p_{A,B}(D)$ solves

$$\frac{1}{1+p} A + \frac{p}{1+p} \{\underline{z}\} \sim_D \frac{1}{1+p} B + \frac{p}{1+p} Z.$$

[28]One could, instead, consider any measurable function $g: I \to \mathcal{A}$ as an opportunity act. Analyzing choice on this larger domain is technically more involved. In particular, it requires a strengthening of Axiom 3 (Continuity). I focus on the smaller domain of simple opportunity acts, as it is sufficient to identify beliefs over tastes with arbitrary support.

(ii) For $B \succ_D A$, $p_{A,B}(D) = -p_{B,A}(D)$.

Call $p_{A,B}(D)$ the *cost of choosing from B instead of A under event D*.

If $\succ$ can be represented by a CPF representation, then the prize $\underline{z}$ must exist because Z is finite and because $\succ_D$ must obviously satisfy Monotonicity. Note that $p_{A,B}(D) \neq 0$ implies that D is nonnull. Endow $\mathcal{A}$ with the topology generated by the Hausdorff metric

$$d_h(A, B) = \max\Big\{\max_{\alpha\in A}\min_{\beta\in B} \|\alpha - \beta\|, \max_{\beta\in B}\min_{\alpha\in A} \|\alpha - \beta\|\Big\}.$$

If two sequences of menus, $\langle A_n\rangle$ and $\langle B_n\rangle$, converge to the same limit in the Hausdorff topology, then the cost of choosing from B_n instead of A_n vanishes under every event. However, the ratio of such costs may have a well defined limit.

Axiom 1′ (Relevance and tightness of objective states). *If $\langle A_n\rangle, \langle B_n\rangle, \langle C_n\rangle \subseteq \mathcal{A}$ converge in the Hausdorff topology, then*

$$\frac{p_{C_n,A_n\cup B_n}(D)}{p_{C_n,B_n}(D)} \nrightarrow 1$$

for some $D \in \mathcal{F}$ implies that there is $D' \in \mathcal{F}$ such that

$$\frac{p_{C_n,A_n}(D')}{p_{C_n,B_n}(D')} \nrightarrow 1.$$

Axiom 1′ implies Axiom 1, where i is substituted by D. To see this, note that Axiom 1 holds trivially unless there is $D \in \mathcal{F}$ such that $A \cup B \nsim_D B$ and $A \sim_D B$. This implies $p_{C,B}(D) = p_{C,A}(D)$ and $p_{C,A\cup B}(D) \neq p_{C,B}(D)$. Define the constant sequences $A_n := A$ and $B_n := B$, and let $C_n := C \succ_D A$. Then $p_{C_n,A_n\cup B_n}(D)/p_{C_n,B_n}(D) \nrightarrow 1$. Thus, according to Axiom 1′, there is $D' \in \mathcal{F}$ with $p_{C_n,A_n}(D')/p_{C_n,B_n}(D') \nrightarrow 1$. Hence $A \nsim_{D'} B$ and Axiom 1 is satisfied. If $p_{C_n,B_n}(D) \nrightarrow 0$, then Axiom 1 also trivially implies Axiom 1′. Thus, Axiom 1′ is only stronger than Axiom 1 for $p_{C_n,B_n}(D) \to 0$.

Theorem 1′. *If $\succ$ has the CPF representation (ϕ, μ, U), then μ and U are unique given ϕ if and only if $\succ$ satisfies Axiom 1′.*

The discussion of the equivalence between (i) and (ii) in Theorem 1 applies here. The intuition for the proof of this equivalence involves identifying taste $s \in S^*$ via two menus, where one is preferred over the other under taste s, but they generate the same payoff under every other relevant taste. If S is continuous, however, then making a menu less preferred by a finite amount under one taste invariably makes it worse under similar tastes (where tastes are viewed as vectors in $\mathbb{R}^k_+$), too. Therefore, individual tastes can only be identified in the limit where the less preferred and the more preferred menu approach each other. In this limit, the cost of choosing from the less preferred menu instead of the more preferred menu tends to zero. Axiom 1′ allows statements about the limit of the ratio of these costs for two different pairs of menus.

In addition to Axioms 2–6, an axiomatization of the CPF representation requires that “small” events do not matter too much for the ranking of acts.

Definition 10. For $f, g \in G$ and $D \in \mathcal{F}$, let fDg be the act, such that

$$fDg(i) := \begin{cases} f(i) & \text{for } i \in D \\ g(i) & \text{otherwise.} \end{cases}$$

Axiom 8 (Event continuity). *For any three acts $f, g, h \in G$ with $h \succ g$, and any sequence $\{D_t\}$ in $\mathcal{F}$ with $D_{t+1} \subset D_t$ and $\bigcap_t D_t = \varnothing$, there exists T such that $h \succ fD_tg$ for all $t > T$.*

Theorem 2′. *The binary relation $\succ$ satisfies Axioms 2–6 and 8 if and only if it has a CPF representation.*

I do not provide a generalization of Theorem 3 here. It would have to be based on a theory that generalizes AA's results to the case of an infinite objective state space. Fishburn (1970, Section 13.3) provides such a generalization.

Appendix B

After collecting some useful properties of support functions,[29] results are established in the order they appear in the text.

B.1 *Support functions*

Definition 11. Call $\sigma_A : S \to \mathbb{R}$ with $\sigma_A(s) := \max_{\alpha \in A}(\alpha \cdot s)$ the *support function* of A.

Support functions have the properties that

(i) $A \subseteq B$ if and only if $\sigma_A \leq \sigma_B$

(ii) $\sigma_{\lambda A+(1-\lambda)B} = \lambda\sigma_A + (1-\lambda)\sigma_B$ whenever $0 \leq \lambda \leq 1$

(iii) $\sigma_{A\cap B} = \sigma_A \wedge \sigma_B$ and $\sigma_{(A\cup B)} = \sigma_A \vee \sigma_B$

(iv) $\sigma_A = \sigma_{\text{conv}(A)}$, where $\text{conv}(A)$ is the convex hull of A.

Denote by A_σ the maximal menu supported by σ, $A_\sigma = \bigcap_{s\in S}\{\alpha \in \Delta(Z) | \alpha \cdot s \leq \sigma(s)\}$. Let $\overline{\mathcal{A}}$ be the collection of all convex subsets of $\Delta(Z)$. Note that $A \in \overline{\mathcal{A}}$ if and only if A is maximal with respect to some support function. Let $\succ_i$ simultaneously denote the induced ranking of support functions: $\sigma \succ_i \xi$ if and only if $A_\sigma \succ_i A_\xi$.

Lemma 2. *For $\varepsilon \geq 0$ small enough, $\sigma_\varepsilon := \varepsilon$ is a support function.*

Proof. The $k-1$ dimensional hyperplane in $\mathbb{R}^k$ that contains S is $H_S = \{x \in \mathbb{R}^k \mid x \cdot \mathbf{1} = 0\}$. The hyperplane that contains the $k-1$ dimensional simplex of lotteries, $\Delta(Z)$, is $H_{\Delta(Z)} = \{x \in \mathbb{R}^k \mid x \cdot \mathbf{1} = 1\}$. These two hyperplanes are parallel. Choose ε

[29]The introduction of support functions to the analysis of choice over menus is a major contribution of DLR. For a comprehensive treatment of support functions in this context, see Chatterjee and Krishna (2011).

small enough such that the $k-1$ dimensional ball $B_\varepsilon \subset H_{\Delta(Z)}$ with radius ε around the center of the simplex is itself inside the simplex, $B_\varepsilon \subset \Delta(Z)$. Then $\sigma_{B_\varepsilon} \equiv \varepsilon$. In particular, the degenerate menu B_0 that contains only the center of the simplex (the lottery that assigns weight $1/k$ to every prize) has support function 0. □

B.2 *Proof of* Theorem 1

Proof of (i) ⇒ (iii). Let $I^\mu \subseteq I^*$ be a largest (in terms of cardinality) subset of linearly independent elements in $\{\mu_i\}_{i\in I^*}$. Then $\#S^* \geq \#I^\mu$ must trivially hold. It has to be shown that $\#S^* = \#I^\mu$. Suppose, to the contrary, that $\#S^* > \#I^\mu$. The definition of S^* implies that one can find at least $\#I^\mu + 1$ distinct Borel sets with non-empty interior, $\{S_t\}_{t=1}^{\#I^\mu+1}$, such that for all $t \leq \#I^\mu + 1$, there exists $i \in I^\mu$ with $\mu_i(\text{int}(S_t)) > 0$. Since μ_i can have at most countably many atoms, one can further guarantee $\mu_i(\text{Cl}(S_t) \cap \text{Cl}(S_{t'})) = 0$ for all $t, t' \leq \#I^\mu + 1$ and all $i \in I^\mu$.

Up to a constant, the vNM expected utility $U_s(\alpha)$ in Definition 4 can be written as $l(s)(s\cdot\alpha)$ with $l(s) > 0$. Then $\max_{\alpha\in A} U_s(\alpha) = l(s)\sigma_A(s)$. As in the text, $l(s)$ captures the "intensity" of taste s.

Claim 1. *Given S_t, there is ε small enough and a support function ξ_t such that $\xi_t = \varepsilon$ on $S \setminus S_t$, $\xi_t \geq \varepsilon$ on S_t, and $x_t(i) := \int_S l(s)[\xi_t(s) - \varepsilon]\,d\mu_i(s) > 0$ for some $i \in I^*$.*

Proof. Remember that σ_ε supports a ball, B_ε, with radius ε around the center of the simplex. The maximal menu B with $\sigma_B \leq \sigma_\varepsilon$ on $S \setminus S_t$ includes all lotteries with $\alpha \cdot s \leq \varepsilon$ for all $s \in S \setminus S_t$. This implies $\max_{\alpha\in B}(\alpha\cdot s) > \varepsilon$ for all s in the nonempty interior of S_t. Hence, $\sigma_B > \sigma_\varepsilon$ must hold on $\text{int}(S_t)$. Let $\xi_t := \sigma_B$. ◁

For x_t as defined in Claim 1,

$$\sum_{t=1}^{\#I^\mu+1} x_t(i)p_t = 0 \quad \text{for all } i \in I^\mu$$

is a system of $\#I^\mu$ independent linear equations with $\#I^\mu + 1$ variables $\{p_t\}_{t\in\{1,\#I^\mu+1\}}$; therefore, it has a nonzero solution, $\sum|p_t| > 0$. Dividing each p_t by $\sum|p_t|$ yields another solution with $\sum|p_t| = 1$. The convex combination of finitely many menus is well defined, and by property (ii) in the previous subsection, the convex combination of finitely many support functions is too. Thus one can define two support functions,

$$\xi := \sum_{t=1}^{\#I^\mu+1} |p_t|(\mathbf{1}_{p_t>0}\xi_t + \mathbf{1}_{p_t<0}\varepsilon)$$

$$\sigma := \sum_{t=1}^{\#I^\mu+1} |p_t|(\mathbf{1}_{p_t>0}\varepsilon + \mathbf{1}_{p_t<0}\xi_t),$$

where ξ_t and ε are as in Claim 1. Then $\sum_{t=1}^{\#I^\mu+1} x_t(i)p_t = 0$ for all $i \in I^\mu$ immediately implies that $A_\xi \sim_i A_\sigma$ for all $i \in I^\mu$. Since I^μ indexes a largest (in terms of cardinality) subset of linearly independent elements in $\{\mu_i\}_{i\in I^*}$, the same must hold for all $i \in I^*$. For $i \in I \setminus I^*$, $A_\xi \sim_i A_\sigma$ trivially holds. At the same time, $p_{t'} \neq 0$ implies that $A_\xi \cup A_\sigma \succ_i A_\xi$ for some $i \in I^\mu$, which contradicts Axiom 1. □

Proof of (ii) ⇒ (i). It has to be established that Axiom 1 is also necessary. Suppose, to the contrary, that the representation exists with the stated uniqueness, but Axiom 1 is violated. Then there are two menus $A, B \in \mathcal{A}$, such that $A \sim_j B$ for all $j \in I$ and $A \cup B \succ_i B$ for some $i \in I$.

Definition 12. The *cost of choosing from* $B \in \mathcal{A}$ *instead of* $A \in \mathcal{A}$ *under taste* $s \in S$ is

$$c_{A,B}(s) := \max_{\alpha\in A} U_s(\alpha) - \max_{\beta\in B} U_s(\beta).$$

The approximation $A \sim_j B$ for all $j \in I$ implies $\sum_{S^*} c_{A,B}(s)\mu_j(s) = 0$ for all $j \in I$. At the same time, $A \cup B \succ_i B$ implies that $c_{A,B}(s)$ cannot be zero under all tastes, so it must be positive under some tastes and negative under others. For the proof, it is important that it is not constant across tastes. Define $\widehat{\mu}(s|i) := (1 + \eta c_{A,B}(s))\mu_i(s)$, where $\eta \neq 0$ is small enough, such that $1 + \eta c_{A,B}(s) > 0$ for all $s \in S^*$. Accordingly define $\widehat{l}(s) := l(s)/(1 + \eta c_{A,B}(s))$. The value function $\widehat{V}$ that corresponds to $(\phi, \widehat{\mu}, \widehat{U})$ is numerically identical to V and, therefore, $(\phi, \widehat{\mu}, \widehat{U})$ also represents $\succ$. This contradicts the uniqueness statement in Theorem 1(ii). Thus, Axiom 1 is necessary for this uniqueness statement. □

Proof of (iii) ⇒ (ii). Suppose (ϕ, μ, U) and $(\phi, \widehat{\mu}, \widehat{U})$ both represent $\succ$. Given the finiteness of S^*, the argument from the third example in Section 2 trivially generalizes to imply that $\mu_i(s)/\mu_j(s) = \widehat{\mu}_i(s)/\widehat{\mu}_j(s)$ for all $i, j \in I^*$ and $s \in S^*$. In particular, $\widehat{\mu}_i(s) = (\widehat{\mu}_1(s)/\mu_1(s))\mu_i(s)$ for all $i \in I^*$ and $s \in S^*$. Since $\widehat{\mu}_i$ is a probability measure on S^* for all $i \in I^*$, then

$$\sum_{s\in S^*} \frac{\widehat{\mu}_1(s)}{\mu_1(s)}\mu_i(s) = 1 \quad \text{for all } i \in I^*.$$

By (iii), this system has $\#S^*$ linearly independent equations in the $\#S^*$ variables $\{\widehat{\mu}_1(s)/\mu_1(s)\}_{s\in S^*}$. Hence, the obvious solution $\widehat{\mu}_1(s)/\mu_1(s) = 1$ for all $s \in S^*$ is unique. This establishes the uniqueness of μ.

Given the uniqueness of μ, the uniqueness of U up to a common rescaling and the addition of constants follows from the identification result in DLR. □

B.3 *Proof of Theorem 2*

Definition 13. As in the proof of Theorem 1, let $\overline{\mathcal{A}}$ be the collection of all convex subsets of $\Delta(Z)$. Let $\overline{G}$ be the collection of all acts, $g: I \to \overline{\mathcal{A}}$. Call $g \in \overline{G}$ a convex valued act.

The proof proceeds to establish that for every act g, there is a convex valued $\overline{g}$ such that $g(i) \sim_i \overline{g}(i)$ for all $i \in I$ and, thus, by Independence, $g \sim \overline{g}$. Additive separability across I is established for convex valued acts. Thus, the act g can be evaluated by finding the act $\overline{g}$ and calculating its value state by state. Finally, Theorem DLRS provides a representation of $\succ_i$ that allows replacing the value of $\overline{g}(i)$ with the subjective expectation of the value of $g(i)$.

Lemma 3. *The binary relation $\succ$ constrained to $\overline{G}$ satisfies Axioms 2–4 if and only if there is a family of continuous linear functions $\{v_i\}_{i\in I}$, $v_i : \overline{\mathcal{A}} \to \mathbb{R}$, such that $v : \overline{G} \to \mathbb{R}$ with $v(g) = \sum_{i\in I} v_i(g(i))$ represents $\succ$ on $\overline{G}$. Moreover, if there is another family of continuous linear functions $\{v'_i\}_{i\in I}$, $v'_i : \overline{\mathcal{A}} \to \mathbb{R}$, such that $v'(g) = \sum_{i\in I} v'_i(g(i))$ represents $\succ$ on $\overline{G}$, then there are constants $a > 0$ and $\{b_i \mid i \in I\}$ such that $v'_i = b_i + av_i$ for each $i \in I$.*

Proof. The collection of convex valued acts $\overline{G}$ together with the convex combination of acts as a mixture operation is a mixture space. Lemma 3 is an application of the mixture space theorem (Theorem 5.11 in Kreps 1988), where additive separability across I follows from the usual induction argument and the continuity of v_i is a consequence of Axiom 2. □

According to Theorem DLRS, $\succ_i$ can be represented by

$$\widehat{V}_i(A) = \int_S \max_{\alpha\in A}(U_{s,i}(\alpha))\, d\widehat{\mu}_i(s)$$

for all $i \in I^*$, where $U_{s,i}$ is a vNM utility function that represents taste s, that is, for any $i \in I^*$ there is a $\widehat{\mu}_i(s)$-measurable function $\lambda_i : S \to \mathbb{R}_+$ such that $\max_{\alpha\in A} U_{s,i}(\alpha) = \lambda_i(s)\sigma_A(s)$. Defining $\mu_i(s) := \lambda_i(s)\widehat{\mu}_i(s)/\int_S \lambda_i(s)\, d\widehat{\mu}_i(s)$ allows $\succ_i$ to be represented by

$$V_i(A) = \int_S \sigma_A(s)\, d\mu_i(s).$$

Corollary 1. *If $i \in I$ is nonnull, then $V_i(A)$ and $v_i(A)$ agree on $\overline{\mathcal{A}}$ up to positive affine transformations.*

Proof. Evaluating $v(g_i^A)$ implies that v_i represents $\succ_i$ on $\overline{\mathcal{A}}$; v_i is linear. The mixture space theorem states that any other linear representation of $\succ_i$ agrees with v_i up to a positive affine transformation. According to Theorem DLRS, $V_i(A)$ is linear and represents $\succ_i$ on $\mathcal{A}$. □

Consequently, there is an event-dependent, positive scaling factor $\pi(i)$ such that, up to a constant, $v_i(A) = \pi(i)V_i(A)$ for all $A \in \overline{\mathcal{A}}$, where $\pi(i) = 0$ if and only if $i \notin I^*$. For every $g \in G$, define $\overline{g} \in \overline{G}$ such that $\overline{g}(i) := \operatorname{conv}(g(i))$ for all $i \in I^*$. Property (iv) of support functions (Appendix B.1) implies $V_i(\overline{g}(i)) = V_i(g(i))$ for all $i \in I^*$. Independence immediately implies that $\overline{g} \sim g$. Let V' represent $\succ$ on G and let $V' \equiv v$ on $\overline{G}$. Then

$V'(g) = V'(\overline{g}) = \sum_{i\in I} v_i(\overline{g}(i)) = \sum_{i\in I} \pi(i)V_i(\overline{g}(i)) = \sum_{i\in I} \pi(i)V_i(g(i))$. Hence, $g \succ h$ if and only if $\sum_{i\in I} \pi(i)V_i(g(i)) > \sum_{i\in I} \pi(i)V_i(h(i))$. Therefore,

$$V'(g) = \sum_{i\in I} \pi(i)\left[\int_S \sigma_{g(i)}(s)\, d\mu_i(s)\right]$$

represents $\succ$. Since v is unique only up to positive affine transformations, $\pi(i)$ can be normalized to be a probability measure, $\phi(i)$. This establishes the sufficiency statement in Theorem 2. In this particular CPF representation, the nonuniqueness of the representation is exploited to normalize the state-independent utilities, U_s, as suggested in DLRS.

That Axioms 2–6 are necessary for the existence of the representation is straightforward to verify.

B.4 *Proof of Theorem 3 (existence)*

If $\succ$ has a CPF representation, then Axiom 7 implies that there is no preference for flexibility on $\Delta(X^*)$. That is, $A \succ_i B$ implies $A \sim_i A \cup B$ for all $A, B \in \Delta(X^*)$ and for all $i \in I$. To see this, suppose, to the contrary, that there is preference for flexibility on $\Delta(X^*)$, that is, there are menus $A, B \subset \Delta(X^*)$ with $A \cup B \succ_i A$ and $A \sim_i B$ for some $i \in I$. But then Axiom 1 implies that there exists $j \in I$ such that $A \nsim_j B$, which contradicts Axiom 7.

Consider the CPF representation from Theorem 2:

$$\widehat{V}(g) = \sum_{i\in I} \widehat{\phi}(i)\left[\int_S \Big(\max_{\alpha\in g(i)} \widehat{U}_s(\alpha)\Big)\, d\widehat{\mu}_i(s)\right].$$

Fix $s' \in S^*$. The fact that there is no preference for flexibility on $\Delta(X^*)$ implies that for any $s \in S^*$, there is $\lambda(s)$ such that $\widehat{U}_s(x) = \lambda(s)\widehat{U}_{s'}(x)$ on X^*, as otherwise one could easily construct $A, B \subset \Delta(X^*)$ with $A \cup B \succ A$. Let $U_s(\cdot) := \widehat{U}_s(\cdot)/\lambda(s)$ to ensure that indeed $U_s(x)$ is constant across S for all $x \in X^*$. Finally, let $\mu_i(s) := \lambda(s)\widehat{\mu}_i(s)/\int_S \lambda(s)\, d\widehat{\mu}_i(s)$ and $\phi(i) := \widehat{\phi}(i) \int_S \lambda(s)\, d\widehat{\mu}_i(s)/(\sum_{i\in I} \widehat{\phi}(i)[\int_S \lambda(s)\, d\widehat{\mu}_i(s)])$ to ensure that

$$V(g) = \sum_{i\in I} \phi(i)\left[\int_S \Big(\max_{\alpha\in g(i)} U_s(\alpha)\Big)\, d\mu_i(s)\right]$$

agrees with $\widehat{V}$ up to a positive affine transformation. □

B.5 *Proof of Theorem 1′*

I first show that Axiom 1′ is sufficient for the uniqueness statement. The definition of support functions (Definition 11) and all related notation remain relevant here. For notational convenience, I omit the dependence of functions from S to $\mathbb{R}$ on $s \in S$ when there is no risk of confusion.

Consider the uninformative event $I \in \mathcal{F}$. Note that $\nu(S') := \int_{S'} l\, d\mu_I(s)$ is itself a positive measure[30]: It exists for any measurable $S' \subset S$, as it is bounded above by $\int_S l\, d\mu_I(s)$, which is finite because the value of the menu supported by σ_ε in Lemma 2 is $\int_S \sigma_\varepsilon l\, d\mu_I(s) = \varepsilon \int_S l\, d\mu_I(s)$. It is positive, because the intensity of tastes $l: S \to \mathbb{R}^+$ is a strictly positive function.[31]

Lemma 4. *There are support functions ξ and σ, and a number $\alpha > 0$ such that $\mu_I(S') - \int_{S'} \alpha(\xi - \sigma) l\, d\mu_I(s) < \varepsilon$. For $\alpha' > \alpha$, there are also support functions ξ' and σ' such that $\mu_I(S') - \int_{S'} \alpha'(\xi' - \sigma') l\, d\mu_I(s) < \varepsilon$.*

Proof.

Claim 2. *If f is positive and integrable under ν, then for any $\varepsilon > 0$, there is a continuous, bounded, positive function $g: S \to \mathbb{R}$ such that $\int_S |f - g|\, d\nu(s) < \varepsilon$.*

Proof. As f and ν are both weakly positive, $\int_S |f\nu|\, ds$ exists. Thus, for every $\varepsilon > 0$, there exists a continuous function $g: S \to \mathbb{R}$ such that $\int_S |g - f|\, d\nu(s) < \varepsilon$. See, for example, Billingsley (1995, Theorem 17.1). Since f is positive, g can be chosen to be positive. ◁

Note that $1/l: S \to \mathbb{R}_+$ is strictly positive because l is. It is integrable under ν because $\int_S 1/l\, d\nu(s) = \int_S d\mu_I(s) = 1$. Given $\varepsilon > 0$, Claim 2 implies that there is a continuous, bounded, positive function g such that

$$\int_{S'} \left| \frac{1}{l} - g \right| d\nu(s) < \tfrac{1}{2}\varepsilon.$$

Claim 3 (Lemma 1.7.9 in Schneider 1993). *The functions that are the difference of two support functions span a cone that is dense in $C(S)$, the space of continuous functions on S, the unit sphere in $\mathbb{R}^k$.*

Claim 3 implies that for every $\varepsilon > 0$, there are two support functions ξ and σ, and a number $\alpha > 0$ such that

$$\int_{S'} |g - \alpha(\xi - \sigma)|\, d\nu(s) < \tfrac{1}{2}\varepsilon$$

for every measurable set $S' \subseteq S$.

Hence,

$$\begin{aligned}\mu_I(S') - \int_{S'} \alpha(\xi - \sigma) l\, d\mu_I(s) &\leq \int_{S'} \left| \frac{1}{l} - \alpha(\xi - \sigma) \right| l\, d\mu_I(s) \\ &\leq \int_{S'} \left| \frac{1}{l} - g \right| d\nu(s) + \int_{S'} |g - \alpha(\xi - \sigma)|\, d\nu(s) < \varepsilon.\end{aligned}$$

[30] If information is ignored, in the sense that DM only gets to choose between degenerate acts that do not condition on information, then preferences can be represented as in DLRS. The measure ν corresponds to the measure featured in this representation. It is dominated by the measure $\mu(s|I)$, and the Radon–Nikodym derivative of ν with respect to $\mu(\cdot|I)$ evaluated in s is $l(s)$, the intensity of taste s.

[31] The equality $l(s) = 0$ corresponds to the trivial state, which is not part of the CPF representation.

This establishes the first part of the lemma. To show the second part, consider $\alpha' = c\alpha$ with $c > 1$. Let $\sigma' = \sigma$ and $\xi' = (1/c)\xi + (1 - 1/c)\sigma$. The variable ξ' is a convex combination of support functions and, therefore, is a support function, and $\alpha'(\xi' - \sigma') \equiv \alpha(\xi - \sigma)$. This concludes the proof of Lemma 4. □

Suppose (ϕ, μ, U) is a CPF representations of $\succ$. Following Lemma 4, one can define a sequence of support functions $\langle \xi_n \rangle$ and $\langle \sigma_n \rangle$, and a sequence of numbers $\langle \alpha_n \rangle$ with $\alpha_n \to \infty$, such that

$$\mu_I(S') - \int_{S'} \alpha_n(\xi_n - \sigma_n) l \, d\mu_I(s) < \frac{1}{n}$$

for every measurable set $S' \subseteq S$ and for all $n > 0$. In particular, $\int_S (\xi_n - \sigma_n) l \, d\mu_I(s) \to 0$.

Now consider another CPF representation of $\succ$, $(\phi, \widehat{\mu}, \widehat{U})$, and define corresponding sequences $\langle \widehat{\xi}_n \rangle$, $\langle \widehat{\sigma}_n \rangle$, and $\langle \widehat{\alpha}_n \rangle$. Obviously, also $\int_S (\widehat{\xi}_n - \widehat{\sigma}_n) \widehat{l} \, d\widehat{\mu}_I(s) \to 0$. It follows immediately from the uniqueness statements in Theorems 3 and 4 in DLR that μ_D and $\widehat{\mu}_D$ share the same support, and that $l(s)\mu_D(s)$ differs from $\widehat{l}(s) \, d\widehat{\mu}_D(s)$ at most by scaling for any $D \in \mathcal{F}$. In particular, $l(s)\mu_I(s) \propto \widehat{l}(s) \, d\widehat{\mu}_D(s)$ and, therefore, $\int_S (\widehat{\xi}_n - \widehat{\sigma}_n) l \, d\mu_I(s) \to 0$. Continuity of the integral implies that it is possible to choose $\langle \xi_n \rangle$, $\langle \sigma_n \rangle$, $\langle \widehat{\xi}_n \rangle$, and $\langle \widehat{\sigma}_n \rangle$ such that

$$\int_S (\xi_n - \sigma_n) l \, d\mu_I(s) = \int_S (\widehat{\xi}_n - \widehat{\sigma}_n) l \, d\mu_I(s)$$

for all $n > 0$ and, hence, $(1/2)\xi_n + (1/2)\widehat{\sigma}_n \sim_I (1/2)\widehat{\xi}_n + (1/2)\sigma_n$ according to (ϕ, μ, U) for all $n > 0$.

Rewriting $p_{A,B}(D)$ as defined in Definition 9 in terms of support functions yields $p_{A,B}(D) \propto \int_S (\sigma_A - \sigma_B) l \, d\mu_D(s)$. For the remainder of the proof, let A_n, B_n, and C_n be defined such that $\sigma_{A_n} = (1/2)\xi_n + (1/2)\widehat{\sigma}_n$, $\sigma_{B_n} = (1/2)\widehat{\xi}_n + (1/2)\sigma_n$, and $\sigma_{C_n} = (1/2)\sigma_n + (1/2)\widehat{\sigma}_n$.

Claim 4. *We have* $p_{C_n, A_n}(D)/p_{C_n, B_n}(D) \to 1$ *for all* $D \in \mathcal{F}$.

Proof. First note that

$$\frac{p_{C_n,A_n}(D)}{p_{C_n,B_n}(D)} = \frac{\int_S \frac{1}{2}(\sigma_n + \widehat{\sigma}_n - \xi_n - \widehat{\sigma}_n) l \, d\mu_D(s)}{\int_S \frac{1}{2}(\sigma_n + \widehat{\sigma}_n - \widehat{\xi}_n - \sigma_n) l \, d\mu_D(s)} = \frac{\int_S (\xi_n - \sigma_n) l \, d\mu_D(s)}{\int_S (\widehat{\xi}_n - \widehat{\sigma}_n) l \, d\mu_D(s)}.$$

By definition, $\mu_I(S') - \alpha_n \int_{S'} (\xi_n - \sigma_n) l \, d\mu_I(s) < 1/n$ for every measurable set $S' \subseteq S$ and for all $n > 0$ implies that (i) $\lim_{n\to\infty} [\alpha_n \int_S (\xi_n - \sigma_n) l \, d\mu_I(s)] = 1$, because μ is a probability measure, and (ii) $\alpha_n(\xi_n - \sigma_n) l \to 1$ almost everywhere according to $\mu_I(s)$. The same observations can be made for $\langle \widehat{\xi}_n \rangle$, $\langle \widehat{\sigma}_n \rangle$, $\langle \widehat{\alpha}_n \rangle$, and $(\phi, \widehat{\mu}, \widehat{U})$.

For every $D \in \mathcal{F}$, the measure μ_D is dominated by μ_I, and $S' \subseteq S$ is μ_D-measurable if and only if it is μ_I-measurable. Hence,

$$\lim_{n\to\infty}\left[\alpha_n \int_S (\xi_n - \sigma_n) l \, d\mu_D(s)\right] = 1$$

for all $D \in \mathcal{F}$. Analogously

$$\lim_{n\to\infty}\left[\widehat{\alpha}_n \int_S (\widehat{\xi}_n - \widehat{\sigma}_n) \widehat{l} \, d\widehat{\mu}_D(s)\right] = 1$$

for all $D \in \mathcal{F}$. As in the case of finite I, it is easy to verify that the fact that the limits are independent of D is meaningful in terms of $\succ$.[32] That is, since (ϕ, μ, U) and $(\phi, \widehat{\mu}, \widehat{U})$ both represent $\succ$, there is also a sequence of numbers $\langle \beta_n \rangle$ such that $\lim_{n\to\infty}[\beta_n \int_S (\widehat{\xi}_n - \widehat{\sigma}_n) l \, d\mu_D(s)] = 1$ for all $D \in \mathcal{F}$. Since $(1/2)\xi_n + (1/2)\widehat{\sigma}_n \sim_I (1/2)\widehat{\xi}_n + (1/2)\sigma_n$ for all $n > 0$, it must be that $\alpha_n/\beta_n \to 1$. Together with observation (ii) above, this implies that $\int_S (\xi_n - \sigma_n) l \, d\mu_D(s) / \int_S (\widehat{\xi}_n - \widehat{\sigma}_n) l \, d\mu_D(s) \to 1$ for all $D \in \mathcal{F}$. □

Claim 5. *If (ϕ, μ, U) and $(\phi, \widehat{\mu}, \widehat{U})$ are two CPF representations of $\succ$ that are distinct beyond the changes permitted in the uniqueness statement of Theorem 1′, then $p_{C_n, A_n \cup B_n}(I) / p_{C_n, B_n}(I) \nrightarrow 1$.*

Proof. First note that

$$\frac{p_{C_n, A_n\cup B_n}(I)}{p_{C_n,B_n}(I)} = \frac{\int_S \frac{1}{2}(\sigma_n + \widehat{\sigma}_n - \max\{\xi_n + \widehat{\sigma}_n, \widehat{\xi}_n + \sigma_n\}) l \, d\mu_I(s)}{\int_S \frac{1}{2}(\sigma_n + \widehat{\sigma}_n - \widehat{\xi}_n - \sigma_n) l \, d\mu_I(s)}$$

$$= \frac{\int_S \max\{\xi_n - \sigma_n, \widehat{\xi}_n - \widehat{\sigma}_n\} l \, d\mu_I(s)}{\int_S (\widehat{\xi}_n - \widehat{\sigma}_n) l \, d\mu_I(s)}.$$

Second, note that $\lim_{n\to\infty}[\widehat{\alpha}_n \int_{S'} (\widehat{\xi}_n - \widehat{\sigma}_n) l \, d\mu_I(s)] = \int_{S'} (l/\widehat{l}) \, d\mu_I(s)$. Hence, on the one hand, $\lim_{n\to\infty}[\widehat{\alpha}_n \int_S (\widehat{\xi}_n - \widehat{\sigma}_n) l \, d\mu_I(s)] = \int_S (l/\widehat{l}) \, d\mu_I(s)$ and, on the other hand, $\lim_{n\to\infty}[\alpha_n \int_S (\widehat{\xi}_n - \widehat{\sigma}_n) l \, d\mu_I(s)] = 1$, as established in the proof of Claim 4. It follows that $\lim_{n\to\infty}(\widehat{\alpha}_n/\alpha_n) = \int_S (l/\widehat{l}) \, d\mu_I(s)$. Recall that $\mu_D(s)$ and $\widehat{\mu}_D(s)$ share the same support, and that $l(s)\mu_D(s)$ differs from $\widehat{l}(s) \, d\widehat{\mu}_D(s)$ at most by scaling for any $D \in \mathcal{F}$. Therefore, if (ϕ, μ, U) and $(\phi, \widehat{\mu}, \widehat{U})$ are distinct in the sense of the claim, then the corresponding functions l and $\widehat{l}$ have to be distinct in the sense that there is $S' \subset S$ such that $\int_{S'} (l/\widehat{l}) \, d\mu_I(s) \neq \mu_I(S') \int_S (l/\widehat{l}) \, d\mu_I(s)$. Without loss of generality, suppose that $\int_{S'} (l/\widehat{l}) \, d\mu_I(s) > \mu_I(S') \int_S (l/\widehat{l}) \, d\mu_I(s)$. Taking all this together gives

$$\lim_{n\to\infty}\left[\widehat{\alpha}_n \int_{S'} (\widehat{\xi}_n - \widehat{\sigma}_n) l \, d\mu_I(s)\right] = \int_{S'} \frac{l}{\widehat{l}} d\mu_I(s) > \mu_I(S') \int_S \frac{l}{\widehat{l}} d\mu_I(s) = \mu_I(S') \lim_{n\to\infty} \frac{\widehat{\alpha}_n}{\alpha_n}$$

or

$$\lim_{n\to\infty}\left[\alpha_n \int_{S'} (\widehat{\xi}_n - \widehat{\sigma}_n) l \, d\mu_I(s)\right] > \mu_I(S').$$

[32]See observation 3 in the proof that item (i) of Theorem 1 implies item (ii).

Therefore, $\lim_{n\to\infty}[\alpha_n \int_S \max\{\xi_n - \sigma_n, \widehat{\xi}_n - \widehat{\sigma}_n\} l \, d\mu_I(s)] > 1$, which implies

$$\frac{\int_S \max\{\xi_n - \sigma_n, \widehat{\xi}_n - \widehat{\sigma}_n\} l \, d\mu_I(s)}{\int_S (\widehat{\xi}_n - \widehat{\sigma}_n) l \, d\mu_I(s)} \nrightarrow 1. \qquad \square$$

The combination of Claims 4 and 5 provides a direct violation of Axiom 1′. Hence, Axiom 1′ implies that (ϕ, μ, U) is unique in the sense of Theorem 1′.

It remains to show that Axiom 1′ is also necessary. The argument requires only slight changes compared to the finite case: Suppose, to the contrary, that the representation holds with the stated uniqueness, but Axiom 1′ is violated. Then there are sequences $\langle A_n\rangle, \langle B_n\rangle, \langle C_n\rangle \subseteq A$, which converge in the Hausdorff topology, with $p_{C_n, A_n \cup B_n}(D)/p_{C_n, B_n}(D) \nrightarrow 1$ for some $D \in F$ and $p_{C_n, A_n}(D')/p_{C_n, B_n}(D') \to 1$ for all $D' \in F$. That $p_{C_n, A_n}(D')/p_{C_n, B_n}(D') \to 1$ for all $D' \in F$ implies that

$$\frac{\int_S c_{A_n, B_n}(s) \, d\mu_{D'}(s)}{\int_S c_{C_n, B_n}(s) \, d\mu_{D'}(s)} \to 0$$

for all $D' \in F$. That $p_{C_n, A_n \cup B_n}(D)/p_{C_n, B_n}(D) \nrightarrow 1$ implies that there is a set $S' \subseteq S$ with $\mu_D(S') > 0$ and

$$\frac{\int_{S'} c_{A_n, B_n}(s) \, d\mu_D(s)}{\int_S c_{C_n, B_n}(s) \, d\mu_D(s)} \nrightarrow 0.$$

In complete analogy to the finite case, define

$$\widehat{\mu}(s|D) := \left(1 + \eta \frac{c_{A_n, B_n}(s)}{\int_S c_{C_n, B_n}(s) \mu_D(s)}\right) \mu_D(s),$$

where η is small enough such that $1 + \eta(c_{A_n, B_n}(s)/\int_S c_{C_n, B_n}(s)\mu_D(s)) > 0$ for all $s \in S$. Another CPF representation $(\phi, \widehat{\mu}, \widehat{U})$ can then be defined in complete analogy to the finite case. Thus, Axiom 1′ must hold.

B.6 *Proof of Theorem 2′*

In analogy to Definition 13, let $\overline{G}$ be the collection of all simple convex valued acts. Let $G_{\{D_t\}} \subset G$ be the collection of all acts that are measurable with respect to the partition $\{D_t\}$ of I in $\mathcal{F}$.

Lemma 3′. *The binary relation $\succ$ constrained to $\overline{G}$ satisfies Axioms 2–4 if and only if there are continuous linear functions $v_D : \overline{\mathcal{A}} \to \mathbb{R}$, indexed by $D \in \mathcal{F}$, that satisfy the following statements.*

(i) The function $v : \overline{G} \to \mathbb{R}$ with $v(g) = \sum_{t=1}^{T} v_{D_t}(g(D_t))$ for $g \in \overline{G} \cap G_{\{D_t\}}$ represents $\succ$.

(ii) If $\{D_t\}_{t=1}^{T}$ is a partition of I, $\tau \subseteq \{1, \ldots, T\}$, and $D = \bigcup_{t\in\tau} D_t$, then $v_D(A) = \sum_{t\in\tau} v_{D_t}(A)$ for all $A \in \overline{\mathcal{A}}$.

Moreover, another collection of continuous linear functions, $v'_D:\overline{\mathcal{A}}\to\mathbb{R}$, satisfies (i) and (ii) if and only if there are constants $a>0$ and a finitely additive function $b:\mathcal{F}\to\mathbb{R}$ such that $v'_D=b(D)+av_D$ for each $D\in\mathcal{F}$.

Proof. That $v(g)=\sum_{t=1}^{T}v_{D_t}(g(D_t))$ for $g\in\overline{G}\cap G_{\{D_t\}}$ represents $\succ$ constrained to $\overline{G}\cap G_{\{D_t\}}$ is implied by Lemma 3. If the simple act g is constant on each element of $\{D_t\}_{t=1}^{T}$, then it is also constant on each element of a finer partition $\{D'_t\}_{t=1}^{T'}$. Let $\tau\subseteq\{1,\dots,T'\}$ be such that $D_t=\bigcup_{t\in\tau}D'_t$ and let $\sharp\tau$ be the number of elements in τ. The usual induction argument yields

$$\frac{1}{\sharp\tau}(g^*(D_1),\dots,g^*(D_{t-1}),A,g^*(D_{t+1}),\dots,g^*(D_T))+\frac{\sharp\tau-1}{\sharp\tau}g^*$$
$$=\sum_{t\in\tau}\frac{1}{\sharp\tau}(g^*(D'_1),\dots,g^*(D'_{t-1}),A,g^*(D'_{t+1}),\dots,g^*(D'_{T'})),$$

and thus $v_{D_t}(A)=\sum_{t\in\tau}v_{D'_t}(A)$, which is item (ii) of the lemma. This implies that $v(g)=\sum_{t=1}^{T}v_{D_t}(g(D_t))$ for $g\in G\cap G_{\{D_t\}}$ represents $\succ$ on G, which is item (i).

The uniqueness statement follows immediately from the uniqueness in Lemma 3, where finite additivity of b is implied by property (ii). That the representation implies continuity and linearity of v, and thus the axioms, is obvious. □

As in the proof of Lemma 2, let B_0 denote the degenerate menu that contains only the center of the simplex $\Delta(Z)$. Given a collection of functions v'_D as in Lemma 3′, let $b(D):=-v'_D(B_0)$ to find a collection of functions $v_D=b(D)+v'_D$ that satisfy (i), (ii), and $v_D(B_0)=0$ for all D. Next I establish that Axiom 8 implies that the functions v_D are countably additive in D.

Claim 6. *Suppose $\succ$ constrained to $\overline{G}$ satisfies Axioms 2–4 and 8, and that the functions v_D with $v_D(B_0)=0$ for all D satisfy (i) and (ii) in Lemma 3′. For a countable collection of disjoint sets in F, $\{D_t\}_{t\geq1}$, let $D:=\bigcup_{t\geq1}D_t$. Then $v_D(A)=\sum_{t\geq1}v_{D_t}(A)$ for all $A\in\overline{\mathcal{A}}$.*

Proof. Given a set $A\in\overline{\mathcal{A}}$, let $f(i)=A$ and $g(i)=B_0$ for all $i\in I$. Given $\varepsilon>0$, choose an act h such that $\varepsilon>v(h)-v(g)>0$ (this is possible by continuity of the value function). Axiom 8 implies that for any nested sequence $\{D_t\}$ in $\mathcal{F}$ with $\bigcap D_t=\varnothing$, there exists T such that $h\succ fD_tg$ for all $t>T$. It follows immediately from Lemma 3′ that $v_{D_t}(A)-v_{D_t}(B_0)=v_{D_t}(A)<\varepsilon$ for all $t>T$. A symmetrical argument establishes that $-v_{D_t}(A)<\varepsilon$ for all $t>T$ and, hence, $v_{D_t}(A)\to0$, which implies countable additivity as claimed (see, for example, Theorem 3.1.1 in Dudley 2002). □

Corollary 1 in the proof of Theorem 2 still holds, where i is replaced with D. That is, on $\overline{\mathcal{A}}$ and up to a positive affine transformation, v_D agrees with the representation V_D of $\succ_D$ as provided by Theorem DLRS:

$$V_D(A)=\int_S\sigma_A(s)\,d\mu_D(s).$$

In complete analogy to the proof of Theorem 2, it can be established that there is an event-dependent, positive scaling factor $\pi(D)$ such that

$$V(g)=\sum_{t=1}^{T}\pi(D_t)\int_S \sigma_{g(D_t)}(s)\,d\mu_{D_t}(s)$$

for $g\in G_{\{D_t\}}$, where V represents $\succ$, and $\pi(D)=0$ if and only if D is trivial. For $D\in\mathcal{F}$ and $S'\in\mathcal{B}$, define the measure

$$\eta(D\times S'):=\frac{\pi(D)\int_{S'}d\mu_D(s)}{\pi(I)\int_S d\mu_I(s)}.$$

Claim 7. *The measure η is a countably additive probability measure on $\{D\times S'\mid D\in\mathcal{F}, S'\in\mathcal{B}\}$.*

Proof. Countable additivity with respect to $\mathcal{F}$ follows immediately from the expression for V above, when choosing g such that $\sigma_{g(D)}(s)=\varepsilon$ for all $s\in S^*$ and all $D\in\mathcal{F}$, which is possible by Lemma 2. Fixing $D\in\mathcal{F}$, countable additivity with respect to $\mathcal{B}$ follows from Theorem DLRS, which implies that μ_D is a countably additive measure (not necessarily a probability measure). □

Claim 7 states that η is countably additive on $\{D\times S'\mid D\in\mathcal{F}, S'\in\mathcal{B}\}$, which is a semiring of sets. It can, therefore, be extended to a countably additive probability measure on the product σ-algebra $\mathcal{F}\otimes\mathcal{B}$. See, for example, Proposition 3.2.4 in Dudley (2002). Then

$$V(g)=\int_{I\times S}\sigma_{g(i)}(s)\,d\eta(i,s).$$

The measure $\eta(i,s)$ can be decomposed into a countably additive marginal distribution $\phi(i):=\eta(i,S)$ on I and a countably additive conditional distribution $\mu_i(s)$ on S given i. The existence of the conditional distribution, $\mu_i(s)$, is implied by, for example, Theorem 10.2.8 in Dudley (2002), after observing that $(I\times S,\mathcal{F}\otimes\mathcal{B},\eta)$ is a probability space and S is a Polish space with the standard metric on $\mathbb{R}^k$. Theorem 10.2.1 in Dudley (2002) establishes that

$$V(g)=\int_I\int_S\sigma_{g(i)}(s)\,d\mu_i(s)\,d\phi(i),$$

as desired.[33]

This completes the proof of the sufficiency statement in Theorem 2′. That Axioms 1–6 are also necessary for the existence of the representation follows as in the case of finite I. The necessity of Axiom 8 follows immediately from the countable additivity of the measure ϕ.

[33] As in the case of finite I, the nonuniqueness of the representation is exploited to normalize all utility functions as suggested in DLRS.

References

Ahn, David S. and Todd Sarver (2013), "Preference for flexibility and random choice." *Econometrica*, 81, 341–361. [505, 517]

Anscombe, Francis J. and Robert J. Aumann (1963), "A definition of subjective probability." *The Annals of Mathematical Statistics*, 34, 199–205. [503]

Billingsley, Patrick (1995), *Probability and Measure*, third edition. Wiley, New York. [527]

Chatterjee, Kaylan and R. Vijay Krishna (2011), "On preferences with infinitely many subjective states." *Economic Theory*, 46, 85–98. [522]

Dekel, Eddie, Barton L. Lipman, and Aldo Rustichini (1998), "Recent developments in modeling unforeseen contingencies." *European Economic Review*, 42, 523–542. [519]

Dekel, Eddie, Barton L. Lipman, and Aldo Rustichini (2001), "Representing preferences with a unique subjective state space." *Econometrica*, 69, 891–934. [503]

Dekel, Eddie, Barton L. Lipman, and Aldo Rustichini (2009), "Temptation driven preferences." *Review of Economic Studies*, 76, 937–971. [518]

Dekel, Eddie, Barton L. Lipman, Aldo Rustichini, and Todd Sarver (2007), "Representing preferences with a unique subjective state space: Corrigendum." *Econometrica*, 75, 591–600. [503]

Dudley, Richard M. (2002), *Real Analysis and Probability*. Cambridge University Press, Cambridge. [531, 532]

Epstein, Larry G. (2006), "An axiomatic model of non-Bayesian updating." *Review of Economic Studies*, 73, 413–436. [518]

Epstein, Larry G. and Kyoungwon Seo (2009), "Subjective states: A more robust model." *Games and Economic Behavior*, 67, 408–427. [518]

Fishburn, Peter C. (1970), *Utility Theory for Decision Making*. Wiley, New York. [522]

Grant, Simon and Edi Karni (2005), "Why does it matter that beliefs and valuations be correctly represented?" *International Economic Review*, 46, 917–934. [518]

Gul, Faruk and Wolfgang Pesendorfer (2001), "Temptation and self-control." *Econometrica*, 69, 1403–1435. [513]

Karni, Edi (2008), "On optimal insurance in the presence of moral hazard." *Geneva Risk and Insurance Review*, 33, 1–18. [518]

Karni, Edi (2011a), "Subjective probabilities on a state space." *American Economic Journal: Microeconomics*, 3 (4), 172–185. [518]

Karni, Edi (2011b), "A theory of Bayesian decision making with action-dependent subjective probabilities." *Economic Theory*, 48, 125–146. [518]

Kopylov, Igor (2009), "Finite additive utility representations for preferences over menus." *Journal of Economic Theory*, 144, 354–374. [518]

Kreps, David M. (1979), "A representation theorem for 'preference for flexibility'." *Econometrica*, 47, 565–577. [503]

Kreps, David M. (1988), *Notes on the Theory of Choice*. Westview Press, Boulder. [514, 525]

Kreps, David M. (1992), "Static choice in the presence of unforeseen contingencies." In *Economic Analysis of Markets and Games* (Partha Dasgupta, Douglas Gale, Oliver Hart, and Eric Maskin, eds.), 258–281, MIT Press, Cambridge. [519]

Krishna, R. Vijay and Philipp Sadowski (2013), "Dynamic preference for flexibility." Unpublished paper. [517]

Maskin, Eric and Jean Tirole (1999), "Unforeseen contingencies and incomplete contracts." *Review of Economic Studies*, 66, 83–114. [518]

Nehring, Klaus (1996), "Preference for flexibility and freedom of choice in a Savage framework." Unpublished paper, University of California, Davis. [518]

Nehring, Klaus (1999), "Preference for flexibility in a Savage framework." *Econometrica*, 67, 101–119. [512, 518]

Ozdenoren, Emre (2002), "Completing the state space with subjective states." *Journal of Economic Theory*, 105, 531–539. [506, 508, 512, 517]

Savage, Leonard J. (1954), *The Foundations of Statistics*. Wiley, New York. [503, 513, 518, 520]

Schneider, Rolf (1993), *Convex Bodies: The Brunn–Minkowski Theory*. Cambridge University Press, Cambridge. [527]

Shenone, Pablo (2010), "Identifying subjective beliefs in subjective state space models." Unpublished paper. [517]

Tirole, Jean (1999), "Incomplete contracts: Where do we stand?" *Econometrica*, 67, 741–781. [519]

von Neumann, John and Oscar Morgenstern (1944), *Theory of Games and Economic Behavior*. Princeton University Press, Princeton. [503]

Submitted 2010-10-10. Final version accepted 2012-7-10. Available online 2012-7-16.