ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

http://www.econometricsociety.org/

Econometrica, Vol. 82, No. 2 (March, 2014), 655–703

DYNAMIC PREFERENCE FOR FLEXIBILITY

R. VIJAY KRISHNA
Duke University, Durham, NC 27708, U.S.A.

PHILIPP SADOWSKI
Duke University, Durham, NC 27708, U.S.A.

Econometrica, Vol. 82, No. 2 (March, 2014), 655–703

DYNAMIC PREFERENCE FOR FLEXIBILITY

BY R. VIJAY KRISHNA AND PHILIPP SADOWSKI[1]

We consider a decision maker who faces dynamic decision situations that involve intertemporal trade-offs, as in consumption–savings problems, and who experiences taste shocks that are transient contingent on the state of the world. We axiomatize a recursive representation of choice over state contingent infinite horizon consumption problems, where uncertainty about consumption utilities depends on the observable state and the state follows a subjective Markov process. The parameters of the representation are the subjective process that governs the evolution of beliefs over consumption utilities and the discount factor; they are uniquely identified from behavior. We characterize a natural notion of greater preference for flexibility in terms of a dilation of beliefs. An important special case of our representation is a recursive version of the Anscombe–Aumann model with parameters that include a subjective Markov process over states and state-dependent utilities, all of which are uniquely identified.

KEYWORDS: Dynamic choice, taste shocks, continuation strategic rationality, recursive Anscombe–Aumann model.

1. INTRODUCTION

UNCERTAINTY ABOUT FUTURE CONSUMPTION UTILITIES influences how economic agents make decisions. A decision maker (DM) who is uncertain about future consumption utilities, perhaps due to uncertain future risk aversion or other taste shocks, prefers not to commit to a course of future action today and, therefore, has a preference for flexibility. For example, DM might be willing to forfeit current consumption if this allows him to delay a decision about future consumption. While this intuition is inherently dynamic, standard models that accommodate preference for flexibility, most prominently Kreps (1979) and Dekel, Lipman, and Rustichini (2001) (henceforth DLR[2]), are static in the sense that there is no intertemporal trade-off (although Kreps (1979) suggests that an infinite horizon model of preference for flexibility is desirable). In this paper, we consider a dynamic environment, allow DM's taste shocks to have an unobservable transient as well as an observable persistent component, and provide foundations for a recursive representation of DM's preferences that is fully identified.

We consider state contingent infinite horizon consumption problems (S-IHCPs) as the domain of choice. Given the collection of relevant states of the world, S-IHCPs are defined recursively as acts that specify, for each state of the world, a decision problem, which is a menu (i.e., a closed set) of lot-

[1]We would like to thank, without implicating, Haluk Ergin, Sujoy Mukerji, Wolfgang Pesendorfer, Todd Sarver, Norio Takeoka, six anonymous referees, and Philippe Jehiel for helpful comments and suggestions, and Vivek Bhattarcharya, Matt Horne, and Justin Valasek for valuable research assistance. Krishna gratefully acknowledges support from the National Science Foundation (Grant SES-1132193).

[2]A relevant corrigendum is Dekel, Lipman, Rustichini, and Sarver (2007) (henceforth DLRS).

 DOI: 10.3982/ECTA10072

teries that yield a consumption prize in the present period and a new S-IHCP starting in the next period. For a simple example of an S-IHCP that features only degenerate lotteries, suppose DM has fixed income y in every period and current wealth w. The relevant state of the world is the per period price of consumption, $\rho > 0$. In each period, DM can choose to consume an amount $m \in [0, \frac{w+y}{\rho}]$ at cost $m\rho$. This will leave him with wealth $w' = w + y - m\rho$ for the next period. Given this technology, we can formulate the consumption problem he faces recursively as a price contingent collection of feasible pairs of current consumption and a new consumption problem for the next period,

$$f_w(\rho) := \left\{(m, f_{w'}) : m \in \left[0, \frac{w+y}{\rho}\right], w' = w + y - m\rho\right\}.$$

In the case where the state space is degenerate (in the example this corresponds to ρ being constant), our domain reduces to that of Infinite Horizon Consumption Problems (IHCPs) introduced by Gul and Pesendorfer (2004), henceforth GP.

The domain of S-IHCPs is rich and allows for complicated future choice behavior.[3] For example, future choices may depend directly on time or on consumption history. We rule out such dependencies so as to stay as close to the standard model as possible.[4] A decision maker who does not anticipate any taste shocks, contingent on the state of the world, is not averse to committing to a state contingent plan of choice. Following Kreps (1979), we refer to this property as *State Contingent Strategic Rationality*. In contrast, a decision maker who anticipates taste shocks should satisfy *Monotonicity*, that is, he always weakly prefers to retain additional options.

We consider taste shocks that are transient (relevant only for one period) contingent on the state of the world. DM should then be strategically rational with respect to his choice of continuation problem beginning in the next period. This is the content of our main behavioral axiom *State Contingent Continuation Strategic Rationality* (S-CSR).

Conceptually, full state contingent strategic rationality would greatly simplify the analysis of DM's behavior. If it is violated with respect to *all* S-IHCPs because of unobservable taste shocks, then it is desirable to understand the nature of this violation by asking if DM is strategically rational with respect to some smaller class of S-IHCPs, for example, as in S-CSR.

Based on S-CSR, we provide a representation of *dynamic preference for flexibility* (DPF) that is the solution to a Bellman equation and can, therefore, be analyzed using standard dynamic programming techniques. The evolution

[3]We only model the initial choice of a consumption problem, but our representation suggests dynamically consistent future choice.

[4]Dependence of choice on the consumption history is central to the notion of habits. We are independently working on a model of habit formation on the domain of IHCPs.

of uncertainty about future consumption utilities is driven by the evolution of the (observable) state of the world, which in turn follows a subjective Markov process. All the parameters of the representation, which are the transition operator on the state space, state contingent beliefs over consumption utilities, and the discount factor, are uniquely identified. In contrast, the static model of DLR cannot identify beliefs over consumption utilities.

The domain of S-IHCPs accommodates important special cases. If the observable state space is degenerate (so there are no observable shocks), then S-CSR becomes *Continuation Strategic Rationality* (CSR) on the domain of IHCPs. The resulting representation features unobservable transient taste shocks that are independent and identically distributed (i.i.d.). If, instead, the analyst can only observe nontrivial preferences over state contingent infinite horizon consumption streams (which provide no flexibility), rather than preferences over all S-IHCPs, then the model reduces to a fully identified recursive version of the Anscombe–Aumann model with state dependent utilities.

The DPF representation is a generalization of "taste shock" models that are commonly used in applied work. Our axioms spell out the testable implications of such models. In particular, if the model features i.i.d. shocks, then choice should satisfy CSR. To test this assumption, consider an agent who must order inventory for the coming season. For example, a farm may need to order seeds for the coming planting season and face shocks to their taste for seeds due to unmodelled environmental conditions. Suppose the order can either be placed long in advance or shortly before the season starts. If the agent is willing to pay a premium to delay placing the order, then CSR is violated and, hence, the i.i.d. model is misspecified: Taste shocks must then have a persistent component. It is possible that this persistent component can be captured by easily observable environmental conditions, such as rainfalls prior to the start of the season. In that case, choice should satisfy S-CSR, which can be tested in a similar fashion.

Our identification results allow us to characterize a notion of *greater preference for flexibility* in terms of a dilation (which is a multidimensional mean preserving spread) of beliefs over consumption utilities.

Identification is also relevant for model based forecasting, which is a central reason for the use of formal models. It involves estimating a relevant parameter in one context so as to forecast outcomes in another. Model based forecasting, thus, requires that (i) the relevant parameter is uniquely identified and (ii) the modeller is willing to assume that the parameter is meaningful outside the context of the observed data. In our model, beliefs over consumption utilities are uniquely identified. Those beliefs might, for example, forecast future choice frequencies of alternatives from continuation problems, as discussed in Sadowski (2013).

The remainder of the paper is structured as follows. Section 2 reviews related literature. Section 3 lays out the basic framework, introduces the representation, explains the axioms, and derives our main theorem, which establishes that the axioms are the behavioral content of the representation. Section 4 discusses important special cases. Section 5 concerns the empirical content of our theory, where Section 5.1 discusses the lack of identification in the static model of DLR and provides a direct intuition for the identification of the parameters in our dynamic model, Section 5.2 suggests specific choice experiments from which an experimenter could elicit those parameters, and Section 5.3 provides a characterization of *greater preference for flexibility* in terms of the parameters. Section 6 concludes. All proofs, as well as a general representation of preference for flexibility with an infinite prize space, are provided in the Appendices.

2. RELATED LITERATURE

Our work builds on standard axiomatic models of preference for flexibility, which follow Kreps (1979) and investigate choice over *menus* of consumption outcomes. The second seminal paper in this literature, DLR, modifies the domain to consider menus of lotteries over outcomes as objects of choice. This facilitates the interpretation of the subjective states as tastes, where a taste is simply a (twice normalized) von Neumann–Morgenstern (vN-M) ranking over consumption outcomes. While menu choice can capture DM's attitude towards the future, implied choice in these models is actually static, in the sense that there is no opportunity for any intertemporal trade-off. We provide a model of preference for flexibility due to uncertain consumption utilities in the tradition of these earlier papers, where the implied choice is dynamic.

As mentioned before, if the set of observable states is degenerate, then our recursive domain of S-IHCPs reduces to that of IHCPs, first analyzed by GP, who show that this recursive domain is well defined. GP provide a dynamic model of consumption with temptation preferences. We allow, instead, for uncertain utilities. The space of possible utilities on our dynamic domain is infinite dimensional, which complicates the analysis. Higashi, Hyogo, and Takeoka (2009) also consider preference for flexibility on the domain of IHCPs. Their model only accommodates preference for flexibility due to a random discount factor that follows an i.i.d. process, but not due to, say, uncertainty about risk aversion, as explained in Section 4.2. Our model nests theirs in terms of behavior. Takeoka (2006) considers choice between menus of menus of lotteries and derives a three period version of DLR. The model cannot capture intertemporal trade-offs and, consequently, the representation is only jointly identified, as in DLR. Rustichini (2002) considers preferences over lotteries over sets of infinite consumption paths. While an additive representation can be obtained, the domain precludes a recursive value function and does not allow unique identification of the representation.

Static models of preference for flexibility lack identification because the set of possible utilities *is* the subjective state space, which means that utilities are (trivially) state dependent. These models, therefore, suffer from the same lack of identification as the Anscombe and Aumann (1963) model with state-dependent utilities. DLR suggest the introduction of a numeraire good (with state-*in*dependent valuation) to their model, so as to identify beliefs uniquely in the same sense that they are uniquely identified in the Anscombe and Aumann model with state independence.[5] If the objective state space is degenerate, then the source of identification in our representation is similar, where the numeraire arises naturally in the form of continuation problems. In contrast, continuation utilities in the DPF representation are state dependent, and identification relies directly on the recursive structure of the representation (Section 5.1 provides intuition for this result).

Sadowski (2013) identifies beliefs in a situation where preference for flexibility depends on the state of the world by requiring the observable states to contain information that is "sufficiently relevant" for future preferences. Instead of choice between menus, Gul and Pesendorfer (2006) study random choice from menus. Observed choice frequencies naturally correspond to a unique measure over utilities, but the scaling of utilities remains arbitrary in their model. Ahn and Sarver (2013) simultaneously model choice between menus and random choice from menus. They achieved full identification by requiring the beliefs in the representation of choice between menus to correspond to frequencies of choice from menus. Dillenberger, Lleras, Sadowski, and Takeoka (2013) fully identify representations of preference for flexibility that feature uncertain future beliefs, rather than uncertain future tastes.

That a decision maker can be uncertain about future preferences in a dynamic setting is noted by Koopmans (1964), who distinguishes between *once-and-for-all* planning, where the agent selects an action for all possible future contingencies, and *piecemeal* planning, where, in each period, the agent chooses an action for the current period while simultaneously narrowing the set of alternatives for the future. Jones and Ostroy (1984) consider a nonaxiomatic (but dynamic) model of choice, where an agent prefers flexibility due to uncertainty about future utilities. They point out that preference for flexibility, for example, in the form of preference for liquidity, is a pervasive theme in economics, and they relate many instances in macroeconomics and finance where DM's dynamic behavior exhibits preference for liquidity in particular and preference for flexibility in general, as in Goldman (1974). They also discuss a notion of *greater variability of beliefs*, which roughly corresponds to our notion of a dilation of beliefs. A special case of our representation (see Corollary 6) is used in empirical work by Hendel and Nevo (2006), who study the problem of a decision maker who maintains an inventory.

A growing literature argues that variations in preferences over time—in particular, in risk aversion—are central in explaining various market behaviors. Indeed, estimated risk aversion has been suggested as a useful index of market sentiment (see, for example, Bollerslev, Gibson, and Zhou (2011)), and there is evidence that the largest component of changes of the equity risk premium is variation in risk aversion, rather than the quantity of risk (see Smith and

[5]Schenone (2010) formalizes this argument.

Whitelaw (2009)). More generally, variation of risk aversion over time has been considered in representative agent settings to improve our understanding of asset pricing phenomena. For instance, Campbell and Cochrane (1999) identify variations in risk premia that correlate with the fundamentals of the economy as a crucial aspect of many dynamic asset pricing phenomena.[6] Bekaert, Engstrom, and Grenadier (2010) show that stochastic risk aversion that is not driven by, or perfectly correlated with, the fundamentals of the economy can explain a wide range of asset pricing phenomena, as well as fit important features of bond and stock markets simultaneously. Our DPF representation can accommodate both kinds of stochastic risk aversion.

An immediate application of our model is provided in Krishna and Sadowski (2012), who investigate a Lucas tree economy with an investment stage and a representative agent whose preferences satisfy our axioms. They show that greater variation in risk aversion results in greater price volatility. Perhaps more surprisingly, they apply our characterization of greater preference for flexibility to find that this variation also reduces investment in a productive asset with liquidity constraint. This type of prediction could help discipline the use of taste shock models in applied work.

3. MODEL AND REPRESENTATION

In this section, we describe the environment and provide a model of dynamic preference for flexibility with observable persistent and unobservable transient taste shocks.

3.1. *Environment*

For a compact metric space Y, let $\mathcal{P}(Y)$ denote the space of probability measures endowed with the topology of weak convergence, so that $\mathcal{P}(Y)$ is compact and metrizable. Let $\mathcal{K}(Y)$ denote the space of closed subsets of a compact metric space Y, endowed with the Hausdorff metric, which makes $\mathcal{K}(Y)$ a compact metric space. Let $S := \{1, \dots, n\}$ be a finite set of *states of the world*. For any metric space Y, $\mathcal{H}(Y) := Y^S$ is the space of acts from S to Y.

Let M be a finite set of consumption prizes with typical member m. A *state contingent infinite horizon consumption problem (S-IHCP)* is an act that specifies, for each state of the world, an (infinite horizon) consumption problem. Such a consumption problem is a menu (i.e., a set) of lotteries that yield a consumption prize in the present period and a new S-IHCP starting in the next period.

[6] In Campbell and Cochrane (1999), the correlation is generated indirectly via habit forming consumption choices. Empirically, however, it is the correlation itself that is important and not the particular mechanism (such as habit formation) that drives it, as Bekaert, Engstrom, and Grenadier (2010) point out.

Let H be the collection of all S-IHCPs. It can be shown that H is linearly homeomorphic to the space of all acts that take values in $\mathcal{K}(\mathcal{P}(M \times H))$.[7] We denote this linear homeomorphism as $H \simeq \mathcal{H}(\mathcal{K}(\mathcal{P}(M \times H)))$. Typical S-IHCPs are $f, g \in H$; typical elements $x, y, z \in \mathcal{K}(\mathcal{P}(M \times H))$ (also written as $\mathcal{K}$) are menus of lotteries over consumption and S-IHCPs; $p, q \in \mathcal{P}(M \times H)$ are typical lotteries; and p_m and p_h denote the marginal distributions of p on M and H, respectively. Each $f \in H$ can be identified with an act that yields a compact set of probability measures over $M \times H$ in every state.

We take the convex sum of sets to be the Minkowski sum; thus, for $x, y \in \mathcal{K}$ and $\lambda \in [0, 1]$, $\lambda x + (1 - \lambda)y := \{\lambda p + (1 - \lambda)q : p \in x, q \in y\}$. Notice that if $x, y \in \mathcal{K}$, then $\lambda x + (1 - \lambda)y$ is also closed and, hence, is in $\mathcal{K}$. Then, for $f, g \in H$, we define $(\lambda f + (1 - \lambda)g)(s) := \lambda f(s) + (1 - \lambda)g(s) \in \mathcal{K}$ for each s. Clearly, H is convex.

When there is no risk of confusion, we identify prizes and continuation problems with degenerate lotteries and lotteries with singleton menus. For example, we denote the lottery over continuation problems that yields x with certainty by x, and denote the lottery that yields current consumption m and continuation problem f with certainty by (m, f). Similarly, $x \in H$ also refers to the act that gives $x \in \mathcal{K}$ in every state.

We explicitly model choice between consumption problems from an ex ante perspective, before consumption begins. That is, we analyze a binary relation $\succsim \subset H \times H$, which we refer to as a *preference*. We let $\succ$ and $\sim$ denote, respectively, the asymmetric and symmetric parts of $\succsim$. The interpretation is that, following the initial choice of an S-IHCP, the state of the world, $s \in S$, is realized and DM gets to choose an alternative from the corresponding menu.

The recursive domain of S-IHCPs is amenable to analysis by stochastic dynamic programming. Its construction follows the *descriptive approach* of Kreps and Porteus (1978) in that it more closely describes how economic agents act and, as mentioned above, embodies what Koopmans (1964) refers to as *piecemeal planning*: instead of choosing a consumption stream that determines consumption for all time, at each instant the decision maker chooses immediate consumption as well as a set of alternatives for the future.

3.2. *Representation of Dynamic Preference for Flexibility*

We now introduce a representation of dynamic preference for flexibility for $\succsim$. The relevant subjective state space (for each s) is the set of all vN-M utility functions over M that are identified up to an additive constant, $\mathcal{U} := \{u \in \mathbb{R}^M : \sum u_i = 0\}$, endowed with the Borel σ-algebra. Subjective states $u \in \mathcal{U}$ are naturally referred to as consumption utilities.[8] Probability measures on

[7]We describe the construction of H in Appendix B.

[8]Subjective states in *static* models of state-dependent expected utility, such as the model in DLRS, are consumption utilities that are normalized up to the addition of constants (i.e., requir-

$\mathcal{U}$ are referred to as subjective beliefs. To ensure that expected consumption utility under a measure μ on $\mathcal{U}$ is well defined, the expected utility from every prize $m \in M$ must be finite. We say a measure μ is *nice* if it satisfies $\mu u := \int_{\mathcal{U}} u \, \mathrm{d}\mu(u) \in \mathcal{U}$; it is *nontrivial* if $\mu(\{\mathbf{0}\}) \neq 1$.

DEFINITION 1: Let $\mathcal{U}$ be defined as above and let $(\mu_s)_{s \in S}$ be a collection of nice probability measures on $\mathcal{U}$ such that for some $s' \in S$, $\mu_{s'}$ is nontrivial. Let Π represent the transition probabilities for a fully connected Markov process on S[9] and let π denote time-0 beliefs about the states in S. Also, let $\delta \in (0, 1)$. A preference $\succsim$ on H has a representation of *dynamic preference for flexibility* (a DPF representation) $((\mu_s)_{s \in S}, \Pi, \delta)$ if $V_0(\cdot) := \sum_s V(\cdot, s)\pi(s)$ represents $\succsim$, where $V(\cdot, s)$ is defined recursively as

$$V(f, s) = \sum_{s' \in S} \Pi(s, s') \left[\int_{\mathcal{U}} \max_{p \in f(s')} \left[u(p_m) + \delta V(p_h, s') \right] \mathrm{d}\mu_{s'}(u) \right] \tag{3.1}$$

and where π is the unique stationary distribution of Π.

In the representation above, V is linear on H, and $V(p_h, \cdot)$ denotes the linear extension (by continuity) of V from H to $\mathcal{P}(H)$, that is, $V(p_h, \cdot) = \int V(g, \cdot) \, \mathrm{d}p_h(g)$. Notice that $V_0(f)$ takes the form

$$V_0(f) = \sum_s \pi(s) \int_{\mathcal{U}} \max_{p \in f(s)} \left[u(p_m) + \delta V(p_h, s) \right] \mathrm{d}\mu_s(u),$$

which follows from the fact that $\pi(s') = \sum_s \pi(s)\Pi(s, s')$ for all $s' \in S$, because π is the unique stationary distribution of Π.

The Markov process represented by Π captures observable and possibly persistent taste shocks, while the state contingent measures $(\mu_s)_{s \in S}$ account for additional unobservable transient shocks. Initial time-0 beliefs are given by π and correspond to DM's long-run (ergodic) beliefs about the distribution of states. Finally, DM's patience is captured by the time preference parameter $\delta \in (0, 1)$.

The DPF representation suggests that DM is uncertain about the per period consumption utility, which captures how much enjoyment he derives from today's consumption, but not about his patience.

---

ing $\sum_i u_i = 0$ for every utility u) *and* the scaling (e.g., $\sum_i u_i^2 = 1$ for every possible utility u). In contrast, as Section 5.1 discusses in detail, the dynamic context allows us to identify the scaling from behavior. Anticipating this, the space $\mathcal{U}$ contains all possible consumption utilities up to the addition of constants. In particular, $u \in \mathcal{U}$ implies $\lambda u \in \mathcal{U}$ for all $\lambda \geq 0$. Importantly, additive constants have no behavioral implications in the context of a linear aggregator, such as the expected utility representation in Definition 1.

[9]The probability of transitioning to state s' from state s is $\Pi(s, s')$. The Markov process with state space S is fully connected if $\Pi(s, s') > 0$ for all $s, s' \in S$.

PROPOSITION 2: *Each DPF representation* $((\mu_s)_{s\in S}, \Pi, \delta)$ *induces a unique continuous function* $V \in C(H \times S)$ *that satisfies equation* (3.1).

The proof is in Appendix A.1.

3.3. *Axioms*

We now discuss our axioms without reference to the DPF representation, but Theorem 1 in Section 3.4 establishes that the axioms are its behavioral content.

To define the induced ranking of menus contingent on the state, $\succsim_s$, fix $x^* \in \mathcal{K}$ and for any $x \in \mathcal{K}$, consider the act

$$f_s^x(s') := \begin{cases} x, & \text{if } s = s', \\ x^*, & \text{otherwise.} \end{cases}$$

For any $x, y \in \mathcal{K}$, let $x \succsim_s y$ if and only if $f_s^x \succsim f_s^y$. Under our axioms below, $\succsim_s$ turns out to be independent of the particular x^* chosen in the definition.

Our first two axioms on $\succsim$ collect standard requirements.

AXIOM 1—Statewise Nontriviality: *Every state* s *is nonnull in the sense that there exist* $x, y \in \mathcal{F}$ *such that* $x \succ_s y$.

AXIOM 2—Continuous Order: $\succsim$ *is* (a) *complete and transitive and* (b) *continuous in the sense that the sets* $\{f : f \succsim g\}$ *and* $\{f : g \succsim f\}$ *are closed*.

The following axiom is the usual independence axiom as applied to our domain.

AXIOM 3—Independence: $f \succ g$ *implies* $\lambda f + (1-\lambda)h \succ \lambda g + (1-\lambda)h$ *for all* $\lambda \in (0, 1)$ *and* $h \in H$.

We are interested in a decision maker who anticipates preference shocks contingent on the state and, hence, values flexibility.

AXIOM 4—Monotonicity: $x \cup y \succsim_s x$ *for all* $x, y \in \mathcal{K}$ *and* $s \in S$.

This is the central axiom in Kreps (1979). It says that additional alternatives are always weakly beneficial.

The next axiom says that DM does not care about correlations between outcomes in M and H, but only cares about the marginal distributions induced by the lotteries in the menu. In particular, if two lotteries induce the same marginal distributions over M and H, then DM does not value the flexibility of having both lotteries available for choice.

AXIOM 5—Separability: *For* $p, q \in \mathcal{P}(M \times H)$, *if* $p_m = q_m$ *and* $p_h = q_h$, *then* $\{p, q\} \sim_s \{p\}$ *for all* $s \in S$.

Our version of Separability is stronger than the version assumed by GP. Variants of the axiom also appear in Higashi, Hyogo, and Takeoka (2009) and Schenone (2010). It is instructive to consider a very specific example of a potential complementarity ruled out by Separability. Consider two S-IHCPs that correspond to meal plans for consecutive evenings, $f_1, f_2 \in H$. Both plans provide the diner with a degenerate choice the first two nights, and do not restrict the available meals in any other period. The plan f_1 yields, with equal probability, either pizza today and salad tomorrow or vice versa, while f_2 yields either pizza both nights or salad both nights, also with equal probability. It does not seem unreasonable that DM would prefer a more varied diet and so rank $f_1 \succ f_2$. However, both plans clearly generate the same marginal distribution over meals in each period, and so Separability rules this out. We note that virtually all dynamic models used in applications satisfy Separability.

Versions of the next two axioms appear in GP, who provided a more detailed discussion. We are interested in stationary preferences, where the ranking of continuation problems does not depend on time. The recursive nature of the domain allows us to capture this notion via the following axiom, which says that if $f \succsim g$, then f is also better than g as a certain continuation problem.

AXIOM 6—Aggregate Stationarity: $f \succsim g$ *if and only if* $\{(m, f)\} \succsim \{(m, g)\}$ *for all* $m \in M$.

The domain of S-IHCPs is rich enough to describe temporal lotteries, as in Kreps and Porteus (1978). We abstract from any preference for the timing of resolution of uncertainty by imposing the following axiom.

AXIOM 7—Singleton Indifference to Timing: $\{\lambda(m, f) + (1 - \lambda)(m, g)\} \sim_s \{(m, \lambda f + (1 - \lambda)g)\}$ *for all* $\lambda \in [0, 1]$ *and* $s \in S$.

The axiom states that in every state s, DM is indifferent between (i) receiving lottery $\lambda(m, f) + (1 - \lambda)(m, g)$, which yields consumption m with certainty and determines whether the degenerate continuation problem will be f or g (early resolution), and (ii) receiving with certainty consumption m and the continuation menu $\lambda f + (1 - \lambda)g$, where uncertainty only resolves in the following period (late resolution).[10]

[10]In particular, suppose state s' is realized tomorrow. The hypothetical choice of lotteries p from $f(s')$ and q from $g(s')$, when considered before the resolution of the lottery over (m, f) and (m, g) in (i), generates the distribution $\lambda p + (1 - \lambda)q$ over outcomes $(m', f') \in M \times H$. This is the same distribution as that of lottery $\lambda p + (1 - \lambda)q$ from $(\lambda f + (1 - \lambda)g)(s')$ in (ii). The only difference is that the uncertainty in (i) resolves in two stages, while the uncertainty in (ii) resolves entirely in the second stage.

If all uncertainty is captured by the objective state s, then $\succsim_s$ should be strategically rational. In the presence of unobservable shocks, strategic rationality will be violated, even contingent on the state. Our central new axiom, *State Contingent Continuation Strategic Rationality* (S-CSR), requires all persistent shocks that are relevant for future consumption tastes to be commonly observed (i.e., captured by the state of the world), while allowing additional unobserved transient taste shocks.

AXIOM 8—S-CSR: $\{(m, f)\} \succsim_s \{(m, g)\}$ *implies* $\{(m, f), (m, g)\} \sim_s \{(m, f)\}$ *for all* $m \in M$ *and* $s \in S$.

The axiom considers the choice of a menu contingent on s. All menus compared here fix first period consumption in state s at m. The axiom says that, contingent on s, there is no preference for flexibility with respect to continuation problems. The axiom does *not* imply that DM is certain about his consumption utility following the first period. It only implies that he does not expect this uncertainty to be resolved prior to the subsequent period's choice. In particular, the axiom is silent on how DM values the state contingent option of retaining the alternatives from *both* x and y for choice two periods ahead. We interpret the ranking $\{(m, f_s^{x\cup y})\} \succ \{(m, f_s^x)\} \succsim \{(m, f_s^y)\}$, which is not precluded by Axiom 8, as the manifestation of DM's uncertainty about his state contingent consumption utility two periods ahead.

Our last axiom holds if the state space is large enough, so that the current objective state captures all past persistent shocks. In the more familiar context of objective uncertainty, states are routinely assumed to follow a (first order) Markov process, and this is justified by assuming that the state space is sufficiently large.

AXIOM 9—History Independence: $\{(m, f_s^x)\} \succsim_{s'} \{(m, f_s^y)\}$ *implies* $\{(m, f_s^x)\} \succsim_{s''} \{(m, f_s^y)\}$ *for all* $x, y \in \mathcal{K}$, $m \in M$, *and* $s, s', s'' \in S$.

In analogy to the definition of $x \succsim_s y$ as DM's ranking of menu x over y contingent on the first period's state being s, one can interpret $\{(m, f_s^x)\} \succsim_{s'} \{(m, f_s^y)\}$ as DM's ranking of x over y contingent on the first and second period's states being s' and s, respectively. The axiom says that this contingent preference is independent of the first state realization.

3.4. *Representation Theorem*

DEFINITION 3: Two probability measures μ and μ' on $\mathcal{U}$ are *identical up to scaling* if there is $\lambda > 0$ such that $\mu(D) = \mu'(\lambda D)$ for all measurable $D \subset \mathcal{U}$, where $\lambda D := \{\lambda u : u \in D\}$. Two collections of measures $(\mu_s)_{s\in S}$ and $(\mu'_s)_{s\in S}$ are *identical up to a common scaling* if, for each s, μ_s and μ'_s are identical up to a scaling that is independent of s.

Given a DPF representation $((\mu_s)_{s\in S}, \Pi, \delta)$ of $\succsim$, scaling $(\mu_s)_{s\in S}$ corresponds to a scaling of the induced value function V.

THEOREM 1: *The binary relation $\succsim$ satisfies Axioms 1–9 if and only if it has a DPF representation, $((\mu_s)_{s\in S}, \Pi, \delta)$. Moreover, the measures $(\mu_s)_{s\in S}$ are unique up to a common scaling, and Π and δ are unique.*

We begin the proof (Appendix A.1) by establishing a representation that is additively separable over S in a straightforward manner. The construction of the DPF representation then has three main steps. First, we find a DLR-type representation of the induced binary relation $\succsim_s$ on menus of lotteries. Since our lotteries are defined on an infinite dimensional space of outcomes, this requires a generalization of the representation theorem in DLR for infinite dimensional prize spaces. Theorem 3 in Appendix A.1 provides such a representation that features an infinite dimensional state space with a measure that is normal (i.e., is outer and inner regular, in the terminology of Aliprantis and Border (1999)), but only finitely additive, and that further lacks the identification properties of the representation in DLR.

The second step relies heavily on S-CSR in constructing a recursive representation where δ may depend on the observable state s. We establish that our axioms—in particular, S-CSR (Axiom 8) and Separability (Axiom 5)—allow us to confine attention to a subjective state space that consists of instantaneous consumption utilities and so is finite dimensional. The carrier of the measure can then be taken to be finite dimensional. This enables us to show that the measure is tight with respect to a compact class of sets and so can be regarded as a regular, countably additive probability measure. We can, therefore, confine attention to a finite dimensional state space and countably additive measures.

The final step consists of showing that there exists exactly one such representation where patience, δ, is state independent. We show this via an application of the Perron Theorem (Theorem 4 in Appendix D). Note, again, that Theorem 1 only identifies consumption utilities up to the addition of constants. Importantly, because additive constants have no behavioral implications in the context of the DPF representation (which is linear), the uniquely identified measures (μ_s) are independent of the normalization of these additive constants; in our case, $\sum u_i = 0$ for all $u \in \mathcal{U}$. Intuition for the unique identification that highlights the role of recursivity is given in Section 5.1.

The Markov process Π has a unique stationary distribution because it is fully connected. To see why this stationary distribution is the appropriate prior π for period 0 behavior, recall that Aggregate Stationarity (Axiom 6) requires that $f \succsim g$ if and only if $\{(m, f)\} \succsim \{(m, g)\}$. Therefore, $V_0(\cdot)$ and $\sum_s \pi(s)V(\cdot, s)$ should represent the same preferences. Let us compute both terms in turn:

$$V_0(f) := \sum_{s'} \pi(s') \int_{\mathcal{U}} \max_{p\in f(s)} \left[u(p_m) + \delta V(p_h, s')\right] \mathrm{d}\mu_{s'}(u)$$

and

$$\begin{aligned}\sum_s \pi(s)V(f,s) &= \sum_s \pi(s)\sum_{s'} \Pi(s,s') \int_{\mathcal{U}} \max_{p\in f(s)}\big[u(p_m)+\delta V(p_h,s')\big]\,\mathrm{d}\mu_{s'}(u) \\ &= \sum_{s'}\sum_s \pi(s)\Pi(s,s') \int_{\mathcal{U}} \max_{p\in f(s)}\big[u(p_m)+\delta V(p_h,s')\big]\,\mathrm{d}\mu_{s'}(u).\end{aligned}$$

It is clear that if π is a stationary distribution of Π, that is, if $\pi(s') = \sum_s \pi(s)\Pi(s,s')$ for all $s' \in S$, then $V_0(\cdot) = \sum_s \pi(s)V(\cdot,s)$ for every $f \in H$, and, hence, $V_0(\cdot)$ and $\sum_s \pi(s)V(\cdot,s)$ do represent the same preferences. Using Singleton Indifference to Timing (Axiom 7), we show in Proposition 31 in Appendix D.2.4 that this is the only possibility, that is, π *must* be the unique stationary distribution of Π.

The requirement that Π be a fully connected Markov process may seem unnecessarily strong. A natural, weaker requirement would be that the Markov process on S is irreducible.[11] The class of fully connected Markov processes is dense in the class of irreducible processes. The small gain in generality does not seem to warrant imposing a weaker, but more involved, version of Axiom 9.

It is instructive to consider other plausible representations of preferences over IHCPs that are not a special case of the DPF representation. First, suppose that $\succsim$ can be represented by a value function V with the same structure as in Theorem 1, where $S = \{1,2\}$ and $\pi = (\frac{1}{2},\frac{1}{2})$, but, contrary to the theorem, DM's patience is state dependent, $\delta_1 \neq \delta_2$, and $\Pi = [\begin{smallmatrix}1 & 0\\ 0 & 1\end{smallmatrix}]$ is not positive. The interpretation is that today DM is uncertain about tomorrow's observable state, but once he learns the state, he does not expect it to ever change again. In that case, there is no recursive representation with state-independent patience, δ, because δ_s must play the role of the (obviously unique) discount factor, contingent on being in state s forever. The measures μ_1 and μ_2 are not identified in this representation. It is easy to verify that the corresponding preferences satisfy all our axioms except History Independence (Axiom 9).

Second, the agent's taste shocks may be persistent or otherwise correlated, even contingent on the state. For example, risk aversion might vary over time with positive correlation. The resulting preferences would violate S-CSR (Axiom 8). We provide foundations for such a representation in Krishna and Sadowski (2013). Third, the agent's utility may depend on what he consumed in the past, for example, due to habit forming preferences or a preference

[11]The Markov process on S with transition probabilities given by Π is *irreducible* if for each $s, s' \in S$, there exists t such that $\Pi^t(s,s') > 0$. An irreducible process has a unique stationary distribution.

for intertemporal diversity as discussed after Separability (Axiom 5). Such consumption-dependent utility is not separable over time and, hence, is ruled out by Axiom 5.

4. SPECIAL CASES

We now consider important special cases of the DPF representation.

4.1. *The DPF Representation With Constant Beliefs*

In an environment without observable states, the domain of S-IHCPs reduces to the space of IHCPs, $Z \simeq \mathcal{K}(\mathcal{P}(M \times Z))$, and the DPF representation reduces to the pair (μ, δ), which induces a continuous function $V : Z \to \mathbb{R}$ that satisfies

$$V(x) = \int_{\mathcal{U}} \max_{p \in x} \big[u(p_m) + \delta V(p_z)\big] \, \mathrm{d}\mu(u)$$

and represents $\succsim$ on Z. We refer to (μ, δ) as a representation of *dynamic preference for flexibility with constant beliefs*. Notice that the identification result in Theorem 1 does not depend on any sort of "richness" assumption on $(\mu_s)_{s \in S}$, so that μ in the DPF representation with constant beliefs is identified up to scaling.

4.2. *Random Discount Factors*

In the environment of Section 4.1, we call a DPF representation with constant beliefs (μ, δ) of $\succsim$ on Z a *random discount factor representation* if there is a utility function $u \in \mathcal{U}$ such that $\mu(\{\lambda u : \lambda \geq 0\}) = 1$. In that case, we can fix u and label subjective states by λ, and following, say, Cochrane (2005), we can interpret $(\delta \int \lambda \, \mathrm{d}\mu)/\lambda$ as a random discount factor.

COROLLARY 4: *The binary relation $\succsim$ satisfies Axioms 1–9 and $\succsim|_{\mathcal{K}(\mathcal{P}(M \times \{f^*\}))}$ is strategically rational if and only if $\succsim$ has a random discount factor representation. This representation is unique up to a scaling of μ.*

The proof is in Appendix E.

Strategic rationality here applies to consumption menus in $\mathcal{K}(\mathcal{P}(M \times \{f^*\}))$, where the continuation problem is fixed at f^*. Intuitively, there is then no subjective uncertainty about the per period consumption preferences, but only about the intensity of those preferences. It follows immediately from the uniqueness of μ up to scaling that the implied distribution over random discount factors as defined above is unique and independent of the choice of u. Higashi, Hyogo, and Takeoka (2009) also consider a representation where the effective discount factor can be random. While our model nests the behavior generated by their model, the parametrization of their model renders it outside the class of models we consider here.

4.3. *The Recursive Anscombe–Aumann Representation*

One may wonder what happens if the analyst can only observe nontrivial preferences over L, the space of *state contingent infinite horizon consumption streams* (S-IHCSs). This is a domain commonly used in applications. In our framework, a consumption stream $\ell \in L$ is just an S-IHCP where all menus are degenerate.[12] We thus refer to $\ell(s)$ as a lottery, and denote by $\ell_m(s)$ and $\ell_\ell(s)$ its marginals on M and L, respectively. Obviously, this smaller domain provides no opportunity to elicit unobservable taste shocks, but it is possible to consider a preference $\gtrsim$ on L that satisfies the restriction of our axioms to L. We say $\gtrsim$ has a *recursive, state dependent, Anscombe–Aumann* representation (recursive Anscombe–Aumann representation for short), $((u_s)_{s\in S}, \Pi, \delta)$, if $V_0(\cdot) := \sum_s \pi(s)V(\cdot, s)$ represents $\gtrsim$, where $V(\cdot, s)$ is defined recursively as

$$V(\ell, s) = \sum_{s'\in S} \Pi(s, s')\big[u_{s'}\big(\ell_m(s')\big) + \delta V\big(\ell_\ell(s'), s'\big)\big],$$

where $u_s \in \mathcal{U}$ for each s with $u_s \neq \mathbf{0}$ for some $s \in S$, and Π and δ are as defined in the DPF representation.

COROLLARY 5: *Let $\gtrsim$ be a binary relation on L. Then $\gtrsim$ is statewise nontrivial on L, and satisfies Axioms 2, 3, and 5–9 restricted to L if and only if $\gtrsim$ has a recursive Anscombe–Aumann representation $((u_s)_{s\in S}, \Pi, \delta)$. Moreover, this representation is unique up to a common scaling of $(u_s)_{s\in S}$.*

The proof is in Appendix E.

Given Theorem 1, the intuition for the result is simple. If $\gtrsim$ on L is statewise nontrivial and satisfies the restriction of Axioms 2, 3, and 5–9 to L, then there is a unique binary relation $\succsim$ that extends $\gtrsim$ to H, satisfies Axioms 1–9, and also satisfies State Contingent Strategic Rationality (i.e., each $\succsim_s$ is fully strategically rational). But then $\succsim$ on H admits a unique DPF representation, $((\mu_s)_{s\in S}, \Pi, \delta)$, and it follows immediately that this representation satisfies $\mu_s(\{u_s\}) = 1$ for some $u_s \in \mathcal{U}$ for each s. Therefore, $((u_s)_{s\in S}, \Pi, \delta)$ is a recursive Anscombe–Aumann representation of $\gtrsim$ that coincides with $\succsim|_L$.

Now suppose that $((u_s)_{s\in S}, \Pi, \delta)$ and $((u_s^*)_{s\in S}, \Pi^*, \delta^*)$ are two recursive Anscombe–Aumann representations of $\gtrsim$ such that (u_s) and (u_s^*) differ by more than a common scaling. Define state-dependent probability measures μ_s and μ_s^* such that $\mu_s(\{u_s\}) = 1 = \mu_s^*(\{u_s^*\})$. It is immediate that $((\mu_s)_{s\in S}, \Pi, \delta)$ and $((\mu_s^*)_{s\in S}, \Pi^*, \delta^*)$ are two DPF representations of $\succsim$ that differ by more than a common scaling of (μ_s), contradicting the uniqueness statement in Theorem 1. Hence, the uniqueness in the corollary follows.

[12]The space L is defined recursively as $L \simeq \mathcal{H}(\mathcal{P}(M \times L))$. It is easy to see that L is a closed and convex subset of H.

The argument above illustrates the advantage, in general, of the domain H over L: Eliciting preferences over L corresponds to finding unconditional expectations over consumption utilities, while eliciting preferences over H corresponds to finding (date t) *conditional* expectations.

Identification of subjective probabilities on the state space in the standard (static) Anscombe and Aumann (1963) model requires two assumptions. First, the ordinal ranking of lotteries must be independent of the state. This assumption is referred to as the State Independence axiom. Second, the cardinal representation of these rankings (the utility) is normalized so as to be state independent. In contrast, the recursive Anscombe–Aumann representation of Corollary 5 features instantaneous as well as continuation utilities that are state dependent. Recursivity and a state-independent discount factor are sufficient to imply that beliefs (which are generated by a subjective Markov process) and state-dependent utilities are fully identified.

4.4. *The Independent Shock Representation*

We now discuss how the model in Hendel and Nevo (2006) maps into our setting. They consider an agent who maintains an inventory of a particular good. In each period, he must choose to purchase combinations of the good in question (the *inside* good) and a bundle of other goods (the *outside* good) from his budget set. He must further decide whether to replenish his inventory, or to draw it down by consuming more or less of the inside good. The price of the good, which follows a Markov process, naturally affects the feasible choices, which means that the consumer faces an S-IHCP, where the price is the observable state of the world.

In each period, the agent also faces taste shocks that affect his utility from consumption of the inside good. In Hendel and Nevo's specification, these shocks do not depend on the realized price. It is easy to see that their agent's preferences have a DPF representation, with the additional constraint that the utility shocks are drawn from a state-independent distribution, that is, $\mu_s = \mu$ for all $s \in S$, where $\mu u \neq \mathbf{0}$. We refer to this as an *independent shock* representation.

In what follows, let L_0 denote the space of state-independent consumption streams and let $X := \mathcal{K}(\mathcal{P}(M \times L_0))$. Notice that $L_0 \subset L \subset \mathcal{H}(X)$.

COROLLARY 6: *The binary relation* $\succsim$ *satisfies Axioms 2–9*, $\succsim|_{L_0}$ *is nontrivial, and* $\succsim_s|_X = \succsim_{s'}|_X$ *for all* $s, s' \in S$ *if and only if* $\succsim$ *admits an independent shock representation. Moreover, this representation is unique up to a scaling of* μ.

The nontriviality of $\succsim|_{L_0}$ implies $\sum_s \pi(s)\mu_s u \neq \mathbf{0}$, that is, the *average* expected utility is nontrivial. Requiring $\succsim_s|_X = \succsim_{s'}|_X$ implies that (i) the representations of $\succsim_s$ and $\succsim_{s'}$ are identical up to scaling, and (ii) that the value of the continuation stream $\ell \in L_0$ is independent of the state. Combining (i) and (ii), the two representations must be identical with $\mu_s = \mu_{s'}$ for all $s, s' \in S$.

Hendel and Nevo's setting leads the agent to maintain an inventory, which is to say, it gives him an unconditional preference for flexibility. There are two channels for this preference: The feasible set changes with the stochastic price and the consumer faces additional taste shocks. Our identification results mean, for example, that Hendel–Nevo's assumption that consumption shocks are log-normally distributed has testable implications.

5. EMPIRICAL CONTENT

In this section, we discuss the empirical content of our model. Section 5.1 explains the lack of identification in the static model of DLR and provides a direct intuition for the source of identification in the dynamic model. Section 5.2 suggests a detailed algorithm to elicit the identified parameters from choice data and Section 5.3 characterizes behavior in terms of the parameters.

5.1. *Identification in Static versus Dynamic Models*

DLR considered preferences over static consumption problems without observable states, that is, over $\mathcal{K}(\mathcal{P}(M))$ with typical element a. The relevant version of DLR's result is that if $\succsim^*_M$ on $\mathcal{K}(\mathcal{P}(M))$ is monotone, continuous, and satisfies Independence, then there exists a probability measure μ^* on $\mathcal{U}$ such that the preference functional

$$U^*_M(a) := \int_{\mathcal{U}} \max_{\alpha\in a} u(\alpha)\,\mathrm{d}\mu^*(u)$$

represents $\succsim^*_M$. In this case, we say that μ^* *represents* $\succsim^*_M$. Consider an example where μ^* is supported on $\{u^1, u^2\}$ with $\mu^*(u^1) = \mu^*(u^2) = \frac{1}{2}$. Suppose also that M represents different levels of monetary prizes and that u^1 corresponds to a state with high risk aversion, while u^2 corresponds to a state with low risk aversion. Clearly, it would be desirable to interpret μ^*, which represents $\succsim^*_M$, as saying that the agent subjectively assesses both high and low risk aversion as being equally likely. Such an interpretation is not possible in the static setting because there are many measures that represent $\succsim^*_M$. To see this, for $i = 1, 2$, let $\lambda_i > 0$ and $\lambda_1 + \lambda_2 = 1$, and define $\tilde{u}^i := u^i/2\lambda_i$ (so that $\tilde{u}^i$ has the same risk aversion as u^i). Also, define the measure $\tilde{\mu}$ on $\{\tilde{u}^1, \tilde{u}^2\}$ as $\tilde{\mu}(\tilde{u}^i) := \lambda_i$. It is easy to see that for any $a \in \mathcal{K}(\mathcal{P}(M))$, we have

$$U^*_M(a) = \sum_{u\in\mathcal{U}} \max_{\alpha\in a} u(\alpha)\mu^*(u) = \sum_{\tilde{u}\in\mathcal{U}} \max_{\alpha\in a} \tilde{u}(\alpha)\tilde{\mu}(\tilde{u}).$$

Since λ_1 and λ_2 are arbitrary, we see that there is a continuum of measures $\tilde{\mu}$ that also represent the same preference $\succsim^*_M$ and for each selection of (λ_1, λ_2), we have different subjective probabilities of either high or low risk aversion. In contrast, in the dynamic setting without observable states as in Section 4.1, $\succsim$

defined over Z has the unique DPF representation with constant beliefs (μ, δ), such that the induced preference, $\succsim_M$, over first period consumption problems, $\mathcal{K}(\mathcal{P}(M))$, is then represented (in the DLR sense) by a unique μ. Therefore, in our dynamic setting, we can interpret μ as a probability measure over, say, different levels of risk aversion.

We now provide some intuition for unique identification of a DPF representation. Let $((\mu_s)_{s\in S}, \Pi, \delta)$ and $((\mu'_s)_{s\in S}, \Pi', \delta')$ be two DPF representations of $\succsim$. Unique identification means that (i) for each $s \in S$, μ_s and μ'_s are identical up to a scaling that is independent of s, (ii) the time preference parameters are identical (i.e., $\delta = \delta'$), and (iii) the objective states evolve according to the same subjective process, namely $\Pi = \Pi'$.

(a) *Uniqueness of μ_s up to scaling*. First, notice that induced preferences over menus in state s are represented by $\int_{\mathcal{U}} \max_{p\in f(s)}[u(p_m) + \delta V(p_h, s)]\, \mathrm{d}\mu_s(u)$. This is an additive expected utility (EU) representation in the fashion of DLR, with the special feature that continuation problems are evaluated independent of u and, hence, provide a natural numeraire, albeit a numeraire that depends on the state $s \in S$. As mentioned in the Introduction, the presence of a numeraire allows the unique identification of the measure μ_s up to scaling, just as state-independent utilities allow the identification of beliefs in the Anscombe–Aumann model. Therefore, μ_s and μ'_s are identical up to scaling, but we have yet to establish that this scaling is independent of s.

(b) *Uniqueness of Π and δ*. For simplicity, let us assume that $\succsim$ restricted to the subdomain L of S-IHCSs is statewise nontrivial. To identify Π and δ, we can then ignore state contingent preference for flexibility by observing only choice over S-IHCSs in L (as defined in Section 4.3). By Corollary 5, the DPF representations $((\mu_s)_{s\in S}, \Pi, \delta)$ and $((\mu'_s)_{s\in S}, \Pi', \delta')$ become recursive Anscombe–Aumann representations $((\bar{u}_s)_{s\in S}, \Pi, \delta)$ and $((\bar{u}'_s)_{s\in S}, \Pi', \delta')$, respectively, where as before, $\bar{u}_s = \mu_s u$ and $\bar{u}'_s = \mu'_s u$. Because V and V' both represent $\succsim$, it is easy to see that $\bar{u}_s$ and $\bar{u}'_s$ must represent the same preference for instantaneous consumption in state s, and so must be collinear. That is, for each s, there exists $\lambda_s \in \mathbb{R}_+$ such that $\bar{u}'_s = \lambda_s \bar{u}_s$. Crucially, recursivity allows us to express $V(\cdot, s)|_L$ in terms of $(\bar{u}_s)_{s\in S}$.

Recall that p^*_m is the lottery such that $u(p^*_m) = 0$ for all $u \in \mathcal{U}$. Fix a lottery $\alpha \neq p^*_m$. Suppose that, contingent on being in state s today, DM is indifferent between (i) receiving consumption lottery α today and p^*_m in all other periods, and (ii) receiving β tomorrow in state s', and p^*_m in all other states and periods.[13] In that case $\bar{u}_s(\alpha) = \delta\Pi(s, s')\bar{u}_{s'}(\beta)$ and analogously for $((\bar{u}'_s)_{s\in S}, \Pi', \delta')$. It is a matter of linear algebra to verify that λ_s must be independent of s. For completeness, this is formally established in Lemma 33 in Appendix F. It then follows immediately that $\Pi = \Pi'$ and $\delta = \delta'$. (Our proof of Theorem 1 does not rely on Lemma 33, but instead invokes the Perron Theorem.)

[13] In the DPF representation, all states are nonnull. There are only finitely many pairs of states and, hence, for α sufficiently close to p^*_m, such β exists for all $s, s' \in S$.

(c) *Uniqueness of* $(\mu_s)_{s\in S}$ *up to common scaling*. We observed above that μ_s is identified up to the same scaling as $V(\cdot, s)$, which, therefore, provides a numeraire in state s that does not depend on u. We have just seen that $\bar{u}_s = \lambda\bar{u}'_s$ for some $\lambda > 0$, $\delta = \delta'$, and $\Pi = \Pi'$ for all $s \in S$. This, in turn, implies that on the subdomain of consumption streams, $V'(\cdot, s) = \lambda V(\cdot, s)$, and, hence, μ_s and μ'_s are identified up to the same scaling $\lambda > 0$ for all $s \in S$.

Recursivity plays two important parts in the argument outlined above. First, for each s, it provides continuation problems as a natural numeraire to imply that each μ_s is unique up to a scaling (that may depend on s). Second, the recursivity and the positivity of the transition matrix tie together all the different numeraires, one for each s, to imply that the measures $(\mu_s)_{s\in S}$ are unique up to a *common* scaling and that Π and δ are unique. We emphasize that this argument depends crucially on the time horizon being infinite. The uniqueness of the identification does not hold in any finite horizon setting for essentially the same reasons that identification fails in the static setting. Roughly put, because we do not have identification in the last period, the choice of representation in the last period is arbitrary, and each such choice then results in a different representation of preference at time 0. The infinite horizon obviates the possibility of such arbitrary choices.

5.2. *Elicitation*

Given the identification result in Theorem 1, one may wonder how the parameters in the representation, $((\mu_s), \Pi, \delta)$, could be elicited experimentally if the experimenter knows (or assumes) that DM satisfies all our axioms and, hence, has a unique DPF representation. As all the parameters have a continuous range, eliciting any parameter exactly from binary choice data requires an infinite amount of data. This is a common problem that experimenters address either by directly eliciting indifference points or by approximating indifference points iteratively.[14] We now discuss how to elicit the parameters of the model via indifference points.

Given the true DPF representation, $((\mu_s), \Pi, \delta)$, the average vN-M utility in state s is $\bar{u}_s := \mu_s u$. Directly eliciting each utility index $\bar{u}_s$ up to scaling (where, as before, we ignore additive constants) is a standard exercise on which we do not elaborate here. Theorem 1 implies that elicitation of the parameters is possible even in knife-edge cases where some, or even all, of the $\bar{u}_s$'s are trivial (as long as at least one μ_s is nontrivial), but we assume here for ease of exposition that every $\bar{u}_s$ is nontrivial. In this case, there is a unique normalized utility function $r_s \in \{r \in \mathcal{U} : \sum_i r_i^2 = 1\}$ such that $\bar{u}_s = \|\bar{u}_s\| r_s$ for each s, where we are not yet cognizant of the scaling $\|\bar{u}_s\|$.

[14] As Mas-Colell (1978) notes (in the context of standard consumer preferences), it is possible to precisely pin down the consumer's preferences with ever-increasing data. (Of course, the data have to increase in a "regular" way.) A similar result can be proved in our setting.

(a) To simultaneously elicit $\|\bar{u}_s\|$ (so that $\bar{u}_s = \|\bar{u}_s\| r_s$ is then known), the transition probabilities $\Pi(s, s')$, and the discount factor δ, we can consider the same indifferences as in step (b) of the identification argument in Section 5.1: Fix a lottery $\alpha \neq p_m^*$. Given s and s', find a lottery β such that, contingent on being in state s today, DM is indifferent between (i) receiving consumption lottery α today and p_m^* in all other periods, and (ii) receiving β tomorrow in state s' and p_m^* in all other states and periods. As before, this indifference implies $\bar{u}_s(\alpha) = \delta \Pi(s, s') \bar{u}_{s'}(\beta)$ and, hence,

$$\Pi(s, s') = \frac{1}{\delta} \frac{\bar{u}_s(p_m)}{\bar{u}_{s'}(p'_m)} = \frac{1}{\delta} \frac{\|\bar{u}_s\|}{\|\bar{u}_{s'}\|} \psi_{s,s'},$$

where $\psi_{s,s'} = r_s(p_m)/r_{s'}(p'_m)$ is known. Since $\Pi(s, \cdot)$ must be a probability measure for all s, we must have $\sum_{s'} \psi_{s,s'}/\|\bar{u}_{s'}\| = \delta/\|\bar{u}_s\|$. Let Ψ be the $S \times S$ matrix $\Psi := [\psi_{s,s'}]$. We then need to solve the problem of finding the left eigenvector $\xi \in \Delta^{S-1}$ of Ψ and the corresponding eigenvalue δ so that $\xi \Psi = \delta \xi$.[15] Upon solving this linear programming problem, set $\|\bar{u}_s\| := 1/\xi_s$ and $\Pi(s, s') := \bar{u}_s(p_m)/\delta \bar{u}_{s'}(p'_m)$ for each $s, s' \in S$.

(b) Eliciting the measure μ_s on $\mathcal{U}$ up to scaling is a little more complicated, as we first have to elicit its support up to scaling. To do so, fix $s \in S$ and let $v_s : \mathcal{P}(M) \to \mathbb{R}$ be defined as $v_s(p_m) := \sum_{s'} \Pi(s, s') \bar{u}_{s'}(p_m)$. Let B_ε be the ε-ball around p_m^*. For any $\xi > 0$ that is sufficiently small, it is easy to determine a lottery $\alpha(\xi) \in \mathcal{P}(M)$ such that $\delta v_s(\alpha(\xi)) = \xi$. Then $x := \{(p_m, \alpha(\xi)) : (p_m \in B_\varepsilon) \wedge (\sqrt{\varepsilon^2 + \xi^2} = \kappa)\}$ with κ small enough features a different maximizer, $q^*(u) := \arg\max_x (u + \delta v_s)$, for every possible taste $u + \delta v_s$ with $u \in \mathcal{U}$.[16] Given the menu x, we ask DM to identify the minimal menu x' (in the sense of set inclusion) such that $x' \subset x$ and $x' \sim_s x$. There is then an obvious bijection that maps the elements of x' to the support of μ_s in $\mathcal{U}$, and the mapping is unique up to scaling.

To simplify our argument, let us assume x', and thus the support of μ_s, is finite. If, contingent on state s, DM is indifferent between (i) receiving extra continuation utility ζ when choosing $q^*(u)$ and (ii) receiving extra continuation utility ζ' when choosing $q^*(u')$, then $\zeta \mu_s(u) = \zeta' \mu_s(u')$,[17] and thus we determine the relative weight $\mu_s(u)/\mu_s(u')$. Fixing u, this can be done for all $u' \neq u$, so that the probability measure μ_s is known up to scaling.

[15]Theorem 1 says that there is a unique $\xi \in \Delta^{S-1}$ that solves this problem, that this $\xi \gg \mathbf{0}$, and that the eigenvalue corresponding to ξ is $\delta \in (0, 1)$.

[16]To see this, let $(\varepsilon, \xi)^*(u) := \arg\max_{\{(\varepsilon,\xi):\sqrt{\varepsilon^2+\xi^2}=\kappa\}} \|u\|\varepsilon + \xi$. There is an obvious bijection between $(\varepsilon, \xi)^*$ and $\|u\|$. Furthermore, for $u, u' \in \mathcal{U}$ with $\|u\| = \|u'\|$ and $u \neq u'$, there is a unique maximizer for any vN-M ranking of consumption prizes, where the maximal consumption value in B_ε is $\|u\|\varepsilon$ for all $u \in \mathcal{U}$.

[17]Formally, using the notation from the previous footnote, $x \cup (q_m^*(u), \alpha(\xi^*(u) + \zeta)) \sim x \cup (q_m^*(u'), \alpha(\xi^*(u') + \zeta'))$, where ζ and ζ' must be small enough not to upset the bijection between the maximizers and the support of μ_s.

5.3. *Behavioral Comparison*

Intuitively, one decision maker has more preference for flexibility than another if he has a stronger preference for menus over singletons in every state. In a manner analogous to the characterization of risk aversion where lotteries are compared to certain amounts of money, characterizing preference for flexibility requires a comparison between S-IHCPs and S-IHCSs in L. This comparison is meaningful only if preferences restricted to L are nontrivial.

AXIOM 10—Consumption Nontriviality: *There exist* $\ell, \ell' \in L$ *such that* $\ell \succ \ell'$.

If $\succsim$ has a DPF representation $((\mu_s)_{s\in S}, \Pi, \delta)$, then by Corollary 5 the restriction of $\succsim$ to L is the recursive Anscombe–Aumann representation $((\mu_s u)_{s\in S}, \Pi, \delta)$. It follows that if $\succsim$ has a DPF representation, then it satisfies Axiom 10 if and only if $\mu_s u \neq \mathbf{0}$ for some $s \in S$ (where $\mathbf{0} \in \mathcal{U}$).[18] The recursivity of the representation along with the fact that Π is fully connected implies that if some $\mu_s u \neq \mathbf{0}$, then for all $s' \in S$, $V(\cdot, s')$ restricted to consumption streams is nontrivial.

DEFINITION 7: $\succsim^*$ has a *greater preference for flexibility* than $\succsim$ if $f \succsim \ell$ implies $f \succsim^* \ell$ for all $\ell \in L$ and $f \in H$.[19]

The comparison in the definition implies that $\succsim$ and $\succsim^*$ have the same ranking over S-IHCSs, that is, $\ell \succsim \ell'$ if and only if $\ell \succsim^* \ell'$. (This is Lemma 34 in Appendix F and also assumes Independence.) We require the following definition to compare beliefs.

DEFINITION 8: Let Q be a Markov kernel[20] from $\mathcal{U}$ to itself. Then Q is a *dilation* if it is expectation preserving, that is, if $\int_{\mathcal{U}} u' Q(u, \mathrm{d}u') = u$ for each $u \in \mathcal{U}$. For probability measures μ and μ^* on $\mathcal{U}$, μ^* is a *dilation* of μ if there exists a dilation Q such that $\mu^* = Q\mu$, that is, $\mu^*(\mathrm{d}u') := \int Q(u, \mathrm{d}u')\mu(\mathrm{d}u)$.

If μ^* is a dilation of μ, then $\mu^* u = \mu u$, because a dilation preserves expectations. To facilitate a comparison of two sets of measures, we shall say that

[18]The "if" part of the claim is easy to see. To see the "only if" part, let μ_s have mean $\mu_s u = \mathbf{0}$ for all $s \in S$, and consider $W \in C(H \times S)$ such that $W(\ell, s) = 0$ for all $\ell \in L$ and $W(f, s) \geq 0$ for all $f \in H$. Then, with $\Phi: C(H \times S) \to C(H \times S)$ given by $\Phi W(f, s) := \sum_{s'\in S} \Pi(s, s')[\int_{\mathcal{U}} \max_{p\in x}[u(p_m) + \delta W(p_h, s)] \, \mathrm{d}\mu_s(u)]$, we have $\Phi W(f, s) \geq 0$ for all $f \in H$, with equality on L. Therefore, the unique fixed point of Φ, namely the value function V representing $\succsim$, must also have this property.

[19]The definition generalizes the definition of *more averse to commitment* in Higashi, Hyogo, and Takeoka (2009) to apply to the domain of S-IHCPs. Their Theorem 4.2 is implied by our Theorem 2.

[20]Let $\mathcal{B}$ be the Borel subsets of $\mathcal{U}$, so that $(\mathcal{U}, \mathcal{B})$ is a measurable space. Then $Q: \mathcal{U} \times \mathcal{B} \to [0, 1]$ is a *Markov kernel* from $(\mathcal{U}, \mathcal{B})$ to itself if (i) for each $u \in \mathcal{U}$, $Q(u, \cdot)$ is a probability measure on $(\mathcal{U}, \mathcal{B})$, and (ii) for each $D \in \mathcal{B}$, $Q(\cdot, D)$ is a measurable function defined on $\mathcal{U}$.

a DPF representation $((\mu_s)_{s\in S}, \Pi, \delta)$ is *canonical* if $\|\mu_s u\|_2 = 1$ for the smallest $s \in S = \{1, \dots, n\}$ such that $\mu_s u \neq \mathbf{0}$. Obviously, $\succsim$ admits a canonical DPF representation if and only if $\succsim$ satisfies Axioms 1–10.

THEOREM 2: *Let $\succsim$ and $\succsim^*$ have canonical DPF representations $((\mu_s)_{s\in S}, \Pi, \delta)$ and $((\mu^*_s)_{s\in S}, \Pi^*, \delta^*)$, respectively. Then $\succsim^*$ has a greater preference for flexibility than $\succsim$ if and only if $\Pi = \Pi^*$, $\delta = \delta^*$, and μ^*_s is a dilation of μ_s for each $s \in S$.*

The proof is in Appendix F.

Intuitively, DM* with preference $\succsim^*$ has a greater preference for flexibility than DM with preference $\succsim$ precisely because he expects more uncertainty to resolve before making a choice, which increases the option value of waiting to make a choice from the menu. Some asset pricing implications of Theorem 2 are explored in Krishna and Sadowski (2012).

The behavioral comparison in terms of beliefs provided by Theorem 2 is possible only because the beliefs (μ_s) are identified up to scaling, and Π and δ are identified uniquely in our dynamic setting. In contrast, the notion of "greater preference for flexibility" proposed in DLR for the static context cannot rely on beliefs because those are not identified in their model. Hence, instead of characterizing whether one DM has a stronger preference for flexibility than another, DLR characterized whether one DM has any preference for flexibility whenever the other does. Neither ranking is complete. While ours can only compare preferences that agree on the ranking of singletons, the ranking in DLR can only compare preferences with representations for which the support of the measure is ordered by set inclusion.

We now compare DM's strength of preference for flexibility *across* states.

PROPOSITION 9: *Suppose $\succsim$ has a DPF representation and satisfies Axiom 10. Then, for $s, s' \in S$, $\succsim_s$ exhibits a greater preference for flexibility than $\succsim_{s'}$ if and only if μ_s is a dilation of $\mu_{s'}$ and $\Pi(s, \cdot) = \Pi(s', \cdot)$.*

Proposition 9 describes how DM trades off flexibility across states. For example, if μ_s is a dilation of $\mu_{s'}$ and $\Pi(s, \cdot) = \Pi(s', \cdot)$, and if $\pi(s) \geq \pi(s')$, then flexibility is more valuable to DM in state s than in state s'. If all states are comparable in the sense of greater preference for flexibility, then $\Pi(s, \cdot)$ is independent of s and objective states evolve according to a process that is i.i.d.

6. CONCLUSION

We provide foundations for a uniquely identified recursive representations of preference for flexibility by relaxing the standard assumption of strategic rationality to accommodate unobservable taste shocks. If taste shocks are transient, strategic rationality with respect to continuation problems should be satisfied unconditionally (CSR). If there are transient and persistent taste shocks, but only transient shocks are unobservable, then strategic rationality with respect to continuation problems should hold contingent on the objective state of the world (S-CSR).

What if persistent taste shocks are also unobservable? Then CSR should be satisfied only contingent on some history of past tastes and, in the simplest case, contingent on the current consumption taste. While this taste is by assumption unobservable, CSR would also have to hold contingent on current consumption *choice* from a large enough menu. In Krishna and Sadowski (2013), we pose an axiom called *Choice Contingent CSR* that captures this idea, and derive a representation that can accommodate unobserved persistent taste shocks by allowing consumption utilities themselves to follow a particular and uniquely identified Markov process. The proof builds on the proof for the DPF representation, where the collection of possible consumption rankings plays the role of the state space.

While we suggest the three versions of CSR above—namely CSR along with State Contingent and Choice Contingent CSR—as a natural starting point when trying to understand how strategic rationality is violated, these could obviously be relaxed further. For example, CSR might only be satisfied contingent on finite histories of consumption choices and states. Modelling the corresponding evolution of beliefs as a Markov process would require a larger subjective state space, as well as an assumption that ensures that preference for flexibility is sufficiently persistent. Beyond that, our proofs of identification only require that the implied Markov process over preferences have a unique stationary distribution. Given the persistence of preference for flexibility, such uniqueness is ensured by Stationarity (via an application of the Perron theorem).

Consequently, we view our work as illustrative of how to achieve a fully identified recursive representation of choice when the assumption of strategic rationality is further weakened.

APPENDICES

In what follows, we also consider the space of menus of consumption lotteries, $\mathcal{K}(\mathcal{P}(M))$, with typical members being a, b, and c. By the recursive nature of H, continuation problems are members of H. Let A, B, and C denote typical elements of the collection of menus of continuation lotteries, $\mathcal{K}(\mathcal{P}(H))$. To ease notational burden, we often write $\mathcal{K}$ for $\mathcal{K}(\mathcal{P}(M \times H))$, $\mathcal{K}_M$ for $\mathcal{K}(\mathcal{P}(M))$, and $\mathcal{K}_H$ for $\mathcal{K}(\mathcal{P}(H))$.

APPENDIX A: ABSTRACT REPRESENTATIONS

We first construct a general representation in the spirit of DLRS where the prize space is infinite. We then study the effect of assuming strategic rationality. We begin by defining and collecting some facts about support functions.

Let Y be a compact metric space. Let $\mathcal{P}(Y)$ be the space of probability measures on Y, endowed with the topology of weak convergence, which makes $\mathcal{P}(Y)$ compact and metrizable, and let $C(Y)$ denote the Banach space of uniformly continuous functions on Y. In what follows, for $\mathfrak{u} \in C(Y)$ and $p \in \mathcal{P}(Y)$, we frequently denote $\int \mathfrak{u}\, \mathrm{d}p$ by $\mathfrak{u}(p)$. Fix $p^* \in \mathcal{P}(Y)$ and let $X := \{\mathfrak{u} \in C(Y) : \mathfrak{u}(p^*) = 0\}$.

Define $\mathfrak{U}_Y := \{\mathfrak{u} \in X : \|\mathfrak{u}\|_\infty = 1\}$ and for any weak* closed, convex $G \subset \mathcal{P}(Y)$, let $\sigma_G : \mathfrak{U}_Y \to \mathbb{R}$ given by $\sigma_G(\mathfrak{u}) := \max_{p \in G} \mathfrak{u}(p)$ be its *support function*. The *extended support function* of such a set G is the unique extension of the support function σ_G to X by positive homogeneity. A function defined on X is sublinear, norm continuous, and positively homogeneous if and only if it is the extended support function of some weak* closed, convex subset of $\mathcal{P}(Y)$ (Theorem 5.102 and Corollary 6.27 of Aliprantis and Border (1999)). We, therefore, call a function $\sigma : \mathfrak{U}_Y \to \mathbb{R}$ a support function if its unique extension to X by positive homogeneity is sublinear and norm continuous.

Support functions have the following duality: For any weak* compact, convex subset G of $\mathrm{aff}(\mathcal{P}(Y))$, $G_{\sigma_G} = G$, where $G_{\sigma_G} := \{p \in \mathrm{aff}(\mathcal{P}(Y)) : \mathfrak{u}(p) \leq \sigma_G(\mathfrak{u}) \text{ for all } \mathfrak{u} \in \mathfrak{U}_Y\}$. Support functions also have the following useful properties[21] for weak* compact, convex subsets G, G' of $\mathcal{P}(Y)$: (i) $G \subset G'$ if and only if $\sigma_G \leq \sigma_{G'}$, (ii) $\sigma_{tG+(1-t)G'} = t\sigma_G + (1-t)\sigma_{G'}$ for all $t \in (0, 1)$, (iii) $\sigma_{G \cap G'} = \sigma_G \wedge \sigma_{G'}$, and (iv) and $\sigma_{\mathrm{conv}(G \cup G')} = \sigma_G \vee \sigma_{G'}$. (Since G and G' are convex, it follows from Lemma 5.14 of Aliprantis and Border (1999) that $\mathrm{conv}(G \cup G')$ is also compact.) Notice also that $\sigma_{\{p^*\}} = \mathbf{0}$.

A.1. *Constructing a General Representation*

Let Y be a compact metric space of prizes, so $\mathcal{P}(Y)$ is a compact metric space. Let $\mathcal{K}(\mathcal{P}(Y))$ denote the space of compact subsets of $\mathcal{P}(Y)$, and let $\mathcal{C}(\mathcal{P}(Y))$ denote the space of closed, convex subsets of $\mathcal{P}(Y)$, both endowed with the Hausdorff metric. Then $\mathcal{C}(\mathcal{P}(Y))$ is a closed subspace of $\mathcal{K}(\mathcal{P}(Y))$. We consider a preference $\succsim$ over $\mathcal{K}(\mathcal{P}(Y))$. For notational ease, we write $\mathcal{K}$ and $\mathcal{C}$ for $\mathcal{K}(\mathcal{P}(Y))$ and $\mathcal{C}(\mathcal{P}(Y))$, respectively. Typical elements of $\mathcal{K}$ are denoted by G, G', and so forth. As always, $\alpha G + (1-\alpha)G' \in \mathcal{K}$ is the Minkowski sum.

The space of all vN-M utility functions is simply $C(Y)$, the Banach space of uniformly continuous functions on Y. In general, in a state-dependent additive EU representation, the vN-M utility functions need only be identified up to positive affine transformation. Fix $p^* \in \mathcal{P}(Y)$ as above and, analogous to DLRS, let the subjective state space be given by $\mathfrak{U}_Y := \{\mathfrak{u} \in C(Y) : \mathfrak{u}(p^*) = 0 \text{ and } \|\mathfrak{u}\|_\infty = 1\}$, the space of all vN-M utility functions that (i) take the value 0 at p^*, (ii) are nontrivial on $\mathcal{P}(Y)$, and (iii) lie on the boundary of the unit ball of $C(Y)$.

[21]See, for instance, page 226 of Aliprantis and Border (1999).

Let $\succsim$ be a preference on $\mathcal{K}$. The axioms to be imposed on $\succsim$ are the same as those imposed on $\succsim_s$, modulo the change in domain. Therefore, for expositional ease, we will simply refer to the corresponding axioms for $\succsim_s$ in the text.

We refer to the state space $\mathfrak{U}_Y$ as the *canonical state space*. Let $\mathcal{A}_{\mathfrak{U}_Y}$ be the Borel algebra of sets in $\mathfrak{U}_Y$, and μ a normal[22] charge on $(\mathfrak{U}_Y, \mathcal{A}_{\mathfrak{U}_Y})$. A pair $(\mathfrak{U}_Y, \mu)$ is a *finitely additive EU representation* of $\succsim$ if $V(x) = \int_{\mathfrak{U}_Y} \max_{p\in x} \mathfrak{u}(p)\, d\mu(\mathfrak{u})$ represents $\succsim$. We say that a finitely additive EU representation $(\mathfrak{U}_Y, \mu)$ is *jointly identified* if, given the state space $\mathfrak{U}_Y$, the charge μ is unique.[23]

THEOREM 3: *A preference*, $\succsim$, *satisfies Nontriviality*, *Continuous Order*, *Monotonicity*, *and Independence if*, *and only if*, *it admits a finitely additive EU representation*.

Note the key differences between Theorem 3 and the additive EU representation theorem of DLR. They establish that given the canonical subjective state space, the measure is unique and countably additive, while the theorem above establishes neither. To see why, note first that our subjective state space is not compact. Therefore, the Riesz Representation Theorem only guarantees that the measure is finitely additive. (Notice that this would remain the case even if one were able to establish uniqueness of the representation.) Second, we are unable to show that the finitely additive EU representation obtained is jointly identified. This is because when the subjective state space is finite dimensional, the span of the space of all support functions is dense in the space of all continuous bounded functions on the subjective state space. We are unable to establish this result in our infinite dimensional setting. Before the formal details are presented, we provide some intuition for the proof.

The proof of the representation naturally extends ideas in DLRS to the infinite dimensional setting. The first step is to show that each menu can be identified (isometrically) with its support function and that support functions live in the space of twice normalized, nontrivial vN-M functions on the prize space Y. (This is exactly as in DLRS.) This allows us to define the subjective state space as a space of all twice normalized, continuous, nonconstant, nontrivial functions on Y. Instead of looking at the space of menus, we can look at the space of support functions, a subset of all continuous bounded functions on the subjective state space.

The second step of the proof shows that any linear functional on the space of menus induces a continuous linear functional on the corresponding space of support functions. Moreover, since the preference $\succsim$ is monotone, this linear

[22] A charge is *outer regular* if every set in $\mathcal{A}_{\mathfrak{U}_Y}$ can be approximated from without by open sets, is *inner regular* if it can be approximated from within by closed sets, and is *normal* if it is both outer and inner regular; see also Definition 10.2 in Aliprantis and Border (1999).

[23] This is the sense in which the representation in DLR is identified.

functional is Lipschitz, and can, therefore, be extended to the space of all continuous bounded functions on the subjective state space. (This step uses the Hahn–Banach Theorem.) The final step uses the Riesz Representation Theorem to show that any linear functional on the space of all continuous bounded functions can be written as an integral with respect to a normal finitely additive measure.

PROOF OF THEOREM 3: By the continuity of $\succsim$, and since $\succsim$ satisfies Independence (Axiom 4), an adaptation of Lemmas 1 and 2 from DLR implies that $G \sim \overline{\text{conv}}(G)$ for each $G \in \mathcal{K}$. (In particular, we have $\overline{\text{conv}}(tG + (1-t)G') = t\,\overline{\text{conv}}(G) + (1-t)\,\overline{\text{conv}}(G')$.) Following DLR, it suffices to restrict attention to $\mathcal{C}$, the space of all weak* compact, convex subsets of $\mathcal{P}(Y)$. Here, $\mathcal{P}(Y)$ is endowed with a metric that induces the weak* topology and $\mathcal{C}$ is endowed with the Hausdorff metric. Let $K_0 := \{\sigma \in C_b(\mathfrak{U}_Y) : \sigma = \sigma_G \text{ for some } G \in \mathcal{C}\}$ denote the affine embedding of $\mathcal{C}$ in $C_b(\mathfrak{U}_Y)$, where $C_b(\mathfrak{U}_Y)$ is the space of bounded and continuous functions on $\mathfrak{U}_Y$, endowed with the supremum norm, and recall that by construction, $\mathbf{0} \in K_0$.

Define the induced preference $\succsim^*$ on K_0, so that $G \succsim G'$ if and only if $\sigma_G \succsim^* \sigma_{G'}$. It is easily seen that $\succsim^*$ is complete and transitive, and satisfies Continuity, Independence, and Monotonicity,[24] since $\succsim$ has these properties. By the Mixture Space Theorem (see, for instance, Fishburn (1970) or Kreps (1988)), there exists a function $\varphi : K_0 \to \mathbb{R}$ that represents $\succsim^*$, is linear in the sense that for all $\sigma^1, \sigma^2 \in K_0$ and $\alpha \in [0,1]$, $\varphi(\alpha\sigma^1 + (1-\alpha)\sigma^2) = \alpha\varphi(\sigma^1) + (1-\alpha)\varphi(\sigma^2)$, and is *positive* so that $\sigma^1 \geq \sigma^2$ implies $\varphi(\sigma^1) \geq \varphi(\sigma^2)$.

Let $K_1 := \bigcup_{r\geq 0} rK_0$ be the cone generated by K_0. It can be shown, following the arguments in Lemma S10 of the supplement to DLRS, that (i) $K_1 - K_1$ is a vector subspace of $C_b(\mathfrak{U}_Y)$, and (ii) for any $\omega \in K_1 - K_1$, there exist $\sigma^1, \sigma^2 \in K_0$ and $r > 0$ such that $\omega = r(\sigma^1 - \sigma^2)$. Extend φ to $K_1 - K_1$ by linearity and let Φ_0 denote this extension. This is achieved by setting $\Phi_0(\omega) := r[\varphi(\sigma^1) - \varphi(\sigma^2)]$, where for a given $\omega \in K_1 - K_1$, $\sigma^1, \sigma^2 \in K_0$ and $r > 0$ are such that $\omega = r(\sigma^1 - \sigma^2)$. The arguments in Lemma S11 of the supplement to DLRS allow us to establish that Φ_0 is positive, that is, $\omega_1 \geq \omega_2$ implies $\Phi_0(\omega_1) \geq \Phi_0(\omega_2)$.

As in DLRS, we claim that Φ_0 on $K_1 - K_1$ is Lipschitz. To see this, notice that for any $\omega \in K_1 - K_1$, we have $\omega \leq \|\omega\|_\infty \mathbf{1}$ (indeed, $|\omega| \leq \|\omega\|_\infty \mathbf{1}$, where $\mathbf{1}$ is the constant function that maps to 1), so that $\Phi_0(\omega) \leq \|\omega\|_\infty \Phi_0(\mathbf{1})$, that is, Φ_0 is Lipschitz with constant $\Phi_0(\mathbf{1})$.

Since Φ_0 is Lipschitz on a vector subspace $K_1 - K_1$, the Hahn–Banach Theorem allows us to extend it to $C_b(\mathfrak{U}_Y)$, with the extension being denoted by Φ. The Riesz representation theorem (Theorem 13.9, Aliprantis and Border (1999)) allows us to represent the linear functional Φ on $C_b(\mathfrak{U}_Y)$ as an

[24] Recall that $G \supset G'$ if and only if $\sigma_G \geq \sigma_{G'}$. Monotonicity of $\succsim$ requires that if $G \supset G'$, then $G \succsim G'$, which implies that we must have $\sigma_G \succsim^* \sigma_{G'}$. Thus, $\succsim^*$ satisfies Monotonicity in the sense that for $\sigma^1, \sigma^2 \in K_0$, $\sigma^1 \geq \sigma^2$ implies $\sigma^1 \succsim^* \sigma^2$.

integral with respect to a normal charge μ on $\mathcal{A}_{\mathfrak{U}_Y}$, the algebra generated by the open sets of $\mathfrak{U}_Y$, as desired. *Q.E.D.*

A.2. *Strategic Rationality*

Let $\succsim$ be a preference on $\mathcal{K}$.

DEFINITION 10: A preference $\succsim$ is *strategically rational* if $G \succsim G'$ implies $G \sim G \cup G'$.

In this section, we show that if $\succsim$ has a finitely additive EU representation $(\mathfrak{U}_Y, \mu)$, it is strategically rational if and only if μ is carried by a singleton. We begin with a lemma.

LEMMA 11: *For any* $\mathfrak{U}_0 \subset \mathfrak{U}_Y$, *the following statements are equivalent.*
(a) $|\mathfrak{U}_0| = 1$.
(b) *For all* $p, q \in \mathcal{P}(Y)$ *and for all* $\mathfrak{u}_1, \mathfrak{u}_2 \in \mathfrak{U}_0$, $\mathfrak{u}_1(p) \geq \mathfrak{u}_1(q)$ *if and only if* $\mathfrak{u}_2(p) \geq \mathfrak{u}_2(q)$.

PROOF: It is clear that (a) implies (b). To see that (b) implies (a), suppose contrary to (a) we have $\mathfrak{u}_1, \mathfrak{u}_2 \in \mathfrak{U}_0$, $\mathfrak{u}_1 \neq \mathfrak{u}_2$. By the definition of $\mathfrak{U}_Y$, there exists $p^* \in \mathcal{P}(Y)$ such that $\mathfrak{u}(p^*) = 0$ for all $\mathfrak{u} \in \mathfrak{U}_Y$. Moreover, $\|\mathfrak{u}\|_\infty = 1$ for all $\mathfrak{u} \in \mathfrak{U}_Y$. Therefore, $\mathfrak{u}_1$ is not a positive affine transformation of $\mathfrak{u}_2$.

At the same time, (b) implies via the expected utility theorem that $\mathfrak{u}_1$ is an affine transformation of $\mathfrak{u}_2$, which is a contradiction. *Q.E.D.*

DEFINITION 12: The *carrier* of the charge μ is the set $\mathfrak{U}_\mu := \bigcap\{N : N$ is closed, $\mu(N^c) = 0\}$.

The carrier of the charge always exists and is clearly well defined. If the charge is also a measure, then the carrier is referred to as the *support* of the measure. Moreover, given the definition of μ, we have $\mathfrak{U}_\mu \subset \mathfrak{U}_Y$. For any $p, q \in \mathcal{P}(Y)$, define $\mathfrak{U}_{p,q} := \{\mathfrak{u} \in \mathfrak{U}_Y : \mathfrak{u}(p) > \mathfrak{u}(q)\}$ and $\mathfrak{U}^\circ_{p,q} := \{\mathfrak{u} \in \mathfrak{U}_Y : \mathfrak{u}(p) = \mathfrak{u}(q)\}$. Notice that $\mathfrak{U}_{p,q}$ is always open and $\mathfrak{U}^\circ_{p,q}$ is always closed, since p is a continuous (linear) functional on $\mathfrak{U}_Y$, which is a closed set.

LEMMA 13: *If* $\succsim$ *on* $\mathcal{K}$ *is strategically rational, then* $\min\{\mu(\mathfrak{U}_{p,q}), \mu(\mathfrak{U}_{q,p})\} = 0$ *for all* $p, q \in \mathcal{P}(Y)$.

PROOF: Suppose to the contrary there exist p, q with $\min\{\mu(\mathfrak{U}_{p,q}), \mu(\mathfrak{U}_{q,p})\} = \mu(\mathfrak{U}_{p,q}) > 0$. It is clear that $\{p, q\} \not\sim \{p\}, \{q\}$, which violates strategic rationality. *Q.E.D.*

LEMMA 14: *If for all* $p, q \in \mathcal{P}(Y)$, $\min\{\mu(\mathfrak{U}_{p,q}), \mu(\mathfrak{U}_{q,p})\} = 0$, *then* $|\mathfrak{U}_\mu| = 1$.

PROOF: If for all $p, q \in \mathcal{P}(Y)$, $\min\{\mu(\mathfrak{U}_{p,q}), \mu(\mathfrak{U}_{q,p})\} = 0$, then either (i) $\mathfrak{U}_\mu \subset \mathfrak{U}_{p,q} \cup \mathfrak{U}^\circ_{p,q}$ or (ii) $\mathfrak{U}_\mu \subset \mathfrak{U}_{q,p} \cup \mathfrak{U}^\circ_{p,q}$, but not both (by the definition of $\mathfrak{U}_\mu$, and since $\mathfrak{U}_{p,q}$ and $\mathfrak{U}_{q,p}$ are open).

In other words, for all $p, q \in \mathcal{P}(Y)$ and for all $\mathfrak{u}_1, \mathfrak{u}_2 \in \mathfrak{U}_Y$, $\mathfrak{u}_1(p) \geq \mathfrak{u}_1(q)$ if and only if $\mathfrak{u}_2(p) \geq \mathfrak{u}_2(q)$. Lemma 11 now implies that $|\mathfrak{U}_\mu| = 1$, as required. *Q.E.D.*

We may now put all this together, as follows.

PROPOSITION 15: *Let $\succsim$ be a preference on $\mathcal{K}(\mathcal{P}(Y))$ that has a finitely additive representation $(\mathfrak{U}_Y, \mu)$. Then the following statements are equivalent*:

(a) *$\succsim$ is strategically rational.*

(b) *The carrier of the charge μ is a singleton, that is, $|\mathfrak{U}_\mu| = 1$.*

PROOF: It is easy to see that (b) implies (a). Lemmas 13 and 14 show that (a) implies (b). *Q.E.D.*

APPENDIX B: THE DOMAIN OF SIHCPS

We now sketch the construction of state contingent infinite horizon consumption problems (S-IHCPs) introduced in Section 3.1. As in the text, M is a finite set of consumption prizes in any period and S is a finite set of states.

Let $H_1 := \mathcal{H}(\mathcal{K}(\mathcal{P}(M)))$ denote the set of acts that give a closed subset of $\mathcal{P}(M)$ in each state $s \in S$. It is a compact metric space when endowed with the (product) Hausdorff metric. For each $t > 1$, inductively define $H_t := \mathcal{H}(\mathcal{K}(\mathcal{P}(M \times H_{t-1})))$, which is also a compact metric space. Thus, each $f_t \in H_t$ is an act that gives a closed set of probability measures over $M \times H_{t-1}$ in each state $s \in S$. This allows us to define the space $H^* := \bigtimes_{t=1}^\infty H_t$, which is also a compact metric space.

Thus, an element in H^* is a collection (f_t), where $f_t \in H_t$ for each $t \geq 1$. Each f_t contains information about some of the $f_\tau \in H_\tau$ for all $\tau < t$, since each $p_t \in f_t(s)$ is a probability measure over $M \times H_{t-1}$, each $p_{t-1} \in f_{t-1}(s)$ is a probability measure over $M \times H_{t-2}$, and so on. A sequence (f_t) is *consistent* if, roughly speaking, the information in f_t about f_{t-2} is in accord with the information in f_{t-1} about f_{t-2}.[25] The space of *all* consistent sequences, denoted by $H \subset H^*$, is the space of S-IHCPs. A simple adaptation of Theorem A1 of GP shows that there is a homeomorphism between H and $\mathcal{H}(\mathcal{K}(\mathcal{P}(M \times H)))$. In fact, it can be shown that this homeomorphism is an *affine* or *linear* homeomorphism, so that we can define a linear preference on H (i.e., a preference that satisfies Independence) and study the naturally induced preference on $\mathcal{H}(\mathcal{K}(\mathcal{P}(M \times H)))$.

[25]See GP for a precise definition of consistency. It is easy to construct examples of sequences in H^* that are *not* consistent.

APPENDIX C: A SEPARABLE REPRESENTATION

In the process of obtaining our representations, we will frequently find it useful to obtain an intermediate representation on one state space, and then transform the representation so it is defined on another state space. The following lemma is an abstract version of this idea. In what follows, we consider the prize space $M \times Y$. If Y is compact, $M \times Y$ is also compact. This allows us to define the canonical state space $\mathfrak{U}_{M\times Y}$ and with a typical state given by $\mathfrak{u} \in \mathfrak{U}_{M\times Y}$.

LEMMA 16—Change of State Space: *Let* $(\mathfrak{U}_{M\times Y}, \mu)$ *be a finitely additive EU representation of* $\succsim$. *A sufficient condition for* $(\mathfrak{U}'_{M\times Y}, \mu')$ *to be another finitely additive EU representation of* $\succsim$ *is that there exist functions* $\Psi : \mathfrak{U}_{M\times Y} \to \mathfrak{U}'_{M\times Y}$ *and* $(\zeta, \xi) : \mathfrak{U}_{M\times Y} \to \mathbb{R}_{++} \times \mathbb{R}$ *such that* Ψ *is a measurable bijection and* (ζ, ξ) *are integrable, and that satisfy the following statements*:

(a) *For all* $\mathfrak{u}'$ *in the image of* Ψ, $\mathfrak{u}'(p) = \zeta(\mathfrak{u})(\Psi\mathfrak{u})(p) + \xi(\mathfrak{u})$ *for all* $p \in \mathcal{P}(M \times Y)$.

(b) *The bijection* Ψ *is measure preserving, that is, for all measurable* $D' \subset \mathfrak{U}'_{M\times Y}$, $\mu'(D') = \mu(\Psi^{-1}D')$, *and for all measurable* $D \subset \mathfrak{U}_{M\times Y}$, $\mu(D) = \mu'(\Psi D)$.

If the finitely additive EU representation $(\mathfrak{U}_{M\times Y}, \mu)$ *is jointly identified* (*as in Appendix* A.1), *then the condition is also necessary*.

The proof of the sufficiency part of the lemma merely amounts to a change of variable and is an instance of the nonuniqueness encountered in DLR. The necessary part of the lemma is also not difficult, and a statement and proof (albeit, in a slightly different setting) can be found in Schenone (2010).

Separability (Axiom 5) says that if $p, q \in \mathcal{P}(M \times Y)$ are such that their marginals are identical (i.e., $p_m = q_m$ and $p_y = q_y$), then $\{p, q\} \sim \{p\}$. This gives us the following lemma.

LEMMA 17: *Let* $(\mathfrak{U}_{M\times Y}, \mu)$ *be a finitely additive representation and satisfy Separability* (*Axiom* 5). *For* p *and* q *that induce the same marginals* (*i.e.*, $p_m = q_m$ *and* $p_y = q_y$), $\mu\{\mathfrak{u} \in \mathfrak{U}_{M\times Y} : \mathfrak{u}(p) > \mathfrak{u}(q)\} = 0$.

PROOF: If the lemma were not true, we would have $V(\{p, q\}) > V(\{q\})$, which contradicts Separability (Axiom 5). *Q.E.D.*

DEFINITION 18: Let $(\mathfrak{U}_{M\times Y}, \mu)$ be a finitely additive EU representation. The representation is *finitely additive, separable* if for each $\mathfrak{u}$ in the carrier of μ, there exist $u(\mathfrak{u}) \in \mathcal{U}$ and $v(\mathfrak{u}) \in C(Y)$ such that $\mathfrak{u}(p) := u(p_m; \mathfrak{u}) + v(p_y; \mathfrak{u})$ for each $p \in \mathcal{P}(M \times Y)$, and if the mapping $\mathfrak{u} \mapsto (u, v)$ is measurable. For ease of notation, we suppress the dependence of u and v on $\mathfrak{u}$.

LEMMA 19: *A finitely additive EU representation satisfies Separability if and only if it is also a finitely additive separable representation as in*

$$V(G)=\int_{\mathfrak{U}_{M\times Y}}\max_{p\in G}\big[u(p_m)+v(p_y)\big]\,\mathrm{d}\mu(\mathfrak{u}). \tag{C.1}$$

PROOF: The "if" part is clear. We now prove the "only if" part. Fix $\tilde{m}\in M$ and $\tilde{y}\in Y$. Since $\frac{1}{2}(m,y)+\frac{1}{2}(\tilde{m},\tilde{y})$ and $\frac{1}{2}(m,\tilde{y})+\frac{1}{2}(\tilde{m},y)$ have the same marginals, Lemma 17 implies that $\mathfrak{u}(\frac{1}{2}(m,y)+\frac{1}{2}(\tilde{m},\tilde{y}))=\mathfrak{u}(\frac{1}{2}(m,\tilde{y})+\frac{1}{2}(\tilde{m},y))$ for all $\mathfrak{u}$ in the carrier of μ. Since each $\mathfrak{u}$ is linear in probabilities, $\mathfrak{u}((m,y))=\mathfrak{u}((m,\tilde{y}))+\mathfrak{u}((\tilde{m},y))-\mathfrak{u}((\tilde{m},\tilde{y}))$ must be satisfied for each individual state $\mathfrak{u}$ in the carrier of μ.[26] Following GP, we define $u(m):=\mathfrak{u}((m,\tilde{y}))$ and $v(y):=\mathfrak{u}((\tilde{m},y))-\mathfrak{u}((\tilde{m},\tilde{y}))$ to find $\mathfrak{u}((m,\tilde{y}))=u(m)+v(y)$. This allows us to write $V(G)=\int_{\mathfrak{U}_{M\times Y}}\max_{p\in G}[u(p_m)+v(p_y)]\,\mathrm{d}\mu(\mathfrak{u})$, which is the desired separable representation. *Q.E.D.*

C.1. *Reduction of the State Space*

Recall that $\mathcal{U}:=\{u\in\mathbb{R}^M:\sum_i u_i=0\}$ is the collection of all vN-M utility functions on M modulo additive constants. Notice that $\mathfrak{U}_M=\{r\in\mathcal{U}:\|r\|_2=1\}$. Since $\mathfrak{U}_M\subset\mathcal{U}$, we must have $\sum_i r_i=0$, so that $r(p_m^*)=0$ for all $r\in\mathfrak{U}_M$. We now use Lemma 16 to show that for any finitely additive separable representation, the state space can be viewed as (a subset of) $\mathfrak{U}_M\times[0,1]\times\mathfrak{U}_Y$, with a typical vN-M utility function of the form $\gamma r(p_m)+(1-\gamma)w(p_y)$.

PROPOSITION 20: *A preference $\succsim$ over the compact subsets of $\mathcal{P}(M\times Y)$ with a finitely additive representation $(\mathfrak{U}_{M\times Y},\mu)$ satisfies Separability (Axiom 5) if and only if there is a change of state space as in Lemma 16 that allows us to write a finitely additive representation of $\succsim$ of the form*

$$\int_{\mathfrak{U}_M\times[0,1]\times\mathfrak{U}_Y}\max_{p\in G}\big[\gamma r(p_m)+(1-\gamma)v(p_y)\big]\,\mathrm{d}\mu'(r,\gamma,v),$$

where μ' is a normal charge on $\mathfrak{U}_M\times[0,1]\times\mathfrak{U}_Y$.

PROOF: The "if" part is immediate, so we only prove the "only if" part of the proposition. Let $(\mathfrak{U}_{M\times Y},\mu)$ be a finitely additive representation. (Recall that Theorem 3 guarantees the existence of a finitely additive representation, though we have not ruled out the possibility that there could be many such representations.) By Lemma 19, every such representation must have the property

[26]In comparison, the weaker separability axiom in GP would only imply this condition for the preference functional V that aggregates over all states.

that for every $\mathfrak{u}$ in the carrier of μ, $\mathfrak{u}(p) = u(p_m) + v(p_y)$, where $u \in \mathcal{U}$ and $v \in C(Y)$ are vN-M functions and the mapping $\mathfrak{u} \mapsto (u, v)$ is measurable.

Every separable utility function $u(p_m) + v(p_y)$ is of the form $\alpha r(p_m) + \beta w(p_y)$, where $r \in \mathfrak{U}_M$, $w \in \mathfrak{U}_Y$, and $(\alpha, \beta) \in \mathbb{R}^2_+ \setminus (0, 0)$. Let $X := (\mathfrak{U}_M \times \mathbb{R}_+) \times (\mathfrak{U}_Y \times \mathbb{R}_+)$ and consider $\mathfrak{U}_{M\times Y} \cap X$. A finitely additive separable representation is one where $\mu(\mathfrak{U}_{M\times Y} \cap X) = 1$.

An even smaller state space is $\mathfrak{U}_M \times [0, 1] \times \mathfrak{U}_Y$, wherein the utility in state (r, γ, v) is $\gamma r(p_m) + (1 - \gamma)w(p_y)$.

Define $\Psi : \mathfrak{U}_{M\times Y} \cap X \to \mathfrak{U}_M \times [0, 1] \times \mathfrak{U}_Y$ as $\Psi : ((r, \alpha), (w, \beta)) \mapsto (r, \alpha/(\alpha + \beta), w)$. It is clear that Ψ is continuous and, hence, measurable. Moreover, Ψ is also a bijection. To see this, suppose there are α, α', β, and β' such that $\alpha/(\alpha + \beta) = \alpha'/(\alpha' + \beta')$. This implies $\alpha' = \rho\alpha$ and $\beta' = \rho\beta$ for some $\rho > 0$, which is impossible since, as mentioned above, $\mathfrak{U}_{M\times Y} \cap X$ only contains functions that are unique up to scaling.

By Lemma 16, we may define the charge $\mu'(D) := \mu(\Psi^{-1}D)$ on $\mathfrak{U}_M \times [0, 1] \times \mathfrak{U}_Y$ and restrict attention to a separable representation on this state space. Notice that since Ψ is continuous, the charge μ' is normal (i.e., both inner and outer regular). *Q.E.D.*

C.2. *Additively Separable Representation*

We now show that in the presence of some additional assumptions, the normal probability charge μ' in Proposition 20 can be replaced by a regular probability measure, leading to an *additive, separable EU representation*. Recall that μ' is a *regular probability measure* if it is (i) regular (i.e., outer regular and tight) and (ii) a probability measure, that is, is countably additive, has $\mu'(\mathfrak{U}_M \times [0, 1] \times \mathfrak{U}_Y) = 1$, and is defined on the Borel σ-algebra of $\mathfrak{U}_M \times [0, 1] \times \mathfrak{U}_Y$.

DEFINITION 21: Let $(\mathfrak{U}_M \times [0, 1] \times \mathfrak{U}_Y, \mu')$ be a finitely additive separable EU representation. The representation is *Y-simple* if the marginal charge of μ' on $\mathfrak{U}_Y$—given by $\int_{\mathfrak{U}_M\times[0,1]\times D} d\mu(r, \gamma, w)$ for any $D \in \mathcal{A}_{\mathfrak{U}_Y}$—has a finite carrier. It is *Y-trivial* if the carrier of the marginal on $\mathfrak{U}_Y$ is a singleton.

LEMMA 22: *Every finitely additive, separable representation that is Y-simple can be extended uniquely to an additive, separable, Y-simple representation*

$$V(G) = \int_{\mathfrak{U}_M\times[0,1]\times\{w_1,\dots,w_I\}} \max_{p\in G}\big[\gamma r(p_m) + (1 - \gamma)w(p_y)\big]\, d\mu(r, \gamma, w). \tag{C.2}$$

PROOF: As μ' is Y-simple, there must exist an $I \in \mathbb{N}$ with $I > 0$ such that the carrier of μ' is a closed subset of $\mathfrak{U}_M \times [0, 1] \times \{w_1, \dots, w_I\}$, where $\{w_1, \dots, w_I\} \subset \mathfrak{U}_Y$. Since $\mathfrak{U}_M \times [0, 1] \times \{w_1, \dots, w_I\}$ is compact, it follows immediately that the carrier of μ' is also compact.

The charge μ' is normal (see footnote 22), so for any Borel (algebra) measurable $A \subset \mathfrak{U}_M \times [0,1] \times \{w_1, \dots, w_I\}$, the inner regularity of μ' implies that $\mu'(A) = \sup\{\mu(C) : C \subset A \text{ and closed}\}$. Since $\mathfrak{U}_M \times [0,1] \times \{w_1, \dots, w_I\}$ is compact, it follows that any closed $C \subset A$ is also compact. Therefore, the charge μ' is "tight" relative to a compact class of sets, namely the collection of all compact subsets of $\mathfrak{U}_M \times [0,1] \times \{w_1, \dots, w_I\}$. By Theorem 9.12 of Aliprantis and Border (1999), μ' is also countably additive (on the algebra of open sets).

By the Carathéodory Extension Procedure Theorem, μ' can be uniquely extended from the algebra of open sets to the Borel σ-algebra of $\mathfrak{U}_M \times [0,1] \times \{w_1, \dots, w_I\}$. (The Carathéodory Theorem is Theorem 9.22 of Aliprantis and Border (1999). The extension of the measure is unique since μ' is a finite measure.) The unique extension of μ' is written as μ.

Finally, Theorem 10.7 of Aliprantis and Border (1999) says that a measure on a Polish space is regular, and because $\mathfrak{U}_M \times [0,1] \times \{w_1, \dots, w_I\}$ is a compact subset of Euclidean space, it follows that μ is regular. *Q.E.D.*

C.3. *Identification of the Representation*

Recall that in the original, abstract EU representation theorem (Theorem 3), we are unable to establish that μ is a regular measure or that the representation is jointly identified. We have identified additional assumptions under which it is possible to show that μ is a regular measure. We now show that under those same assumptions, the representation can be jointly identified.

PROPOSITION 23: *Suppose a continuous preference $\succsim$ has an additive, separable, Y-simple representation as in equation (C.2). Then the representation is jointly identified, that is, the measure μ is unique on the state space $\mathfrak{U}_M \times [0,1] \times \{w_1, \dots, w_I\}$ (where $I > 0$ is a natural number).*

PROOF: Let $\succsim$ have a utility representation V as in (C.2) and let $\mu_0 = \mu$. Suppose there is another regular measure μ_1 (that need not be a probability measure) on $\mathfrak{U}_M \times [0,1] \times \{w_{I+1}, \dots, w_{I+J}\}$ (where $J > 0$) such that

$$V(G) = \int_{\mathfrak{U}_M \times [0,1] \times \{w_{I+1}, \dots, w_{I+J}\}} \max_{p \in G} \bigl[\gamma r(p_m) + (1-\gamma) w(p_y)\bigr] \, \mathrm{d}\mu_1(r, \gamma, w).$$

It is without loss of generality to consider μ_0 and μ_1 as measures on $\mathfrak{U}_M \times [0,1] \times \{w_1, \dots, w_{I+J}\}$. We show that $\mu_1 = \mu_0$.

Let $Y_0 \subset Y$ be a finite set such that for each $i, j \in 1, \dots, I$, where $i \neq j$, there exists $y_{ij} \in Y_0$ such that $w_i(y_{ij}) \neq w_j(y_{ij})$. The requirement is easily satisfied, since w_i is a continuous, nonconstant function on Y and $i \neq j$ implies $w_i \neq w_j$, which in turn implies that the two functions must disagree somewhere.

Now consider the set $B := M \times Y_0$ and the domain $\mathcal{K}(\mathcal{P}(B))$. Then each measure μ_j, $j = 0, 1$, induces the preference functional W_j on $\mathcal{K}(\mathcal{P}(B))$ as

$$W_j(G) = \int_{\mathfrak{U}_M \times [0,1] \times \{w_1, \ldots, w_{I+J}\}} \max_{p \in G}\big[\gamma r(p_m) + (1-\gamma) w(p_y)\big]\, d\mu_j(r, \gamma, w).$$

Define, as in DLR, $\mathfrak{U}_B := \{r \in \mathbb{R}^B : \sum_{b \in B} r_b = 0, \sum_{b \in B} r_b^2 = 1\}$. In that case, $\mathfrak{U}_M \times [0,1] \times \{w_1, \ldots, w_{I+J}\}$ is isomorphic to a subset of $\mathfrak{U}_B$. Indeed, for any vN-M utility function on $M \times Y_0$ of the form $\alpha r(q_m) + (1-\alpha) w(q_y)$, the two normalizations that map the function into $\mathfrak{U}_B$ are continuous.

Thus, μ_0 and μ_1 are transformed into measures on $\mathfrak{U}_B$ in the obvious way and have supports on sets (in $\mathfrak{U}_B$) that are isomorphic to $\mathfrak{U}_M \times [0,1] \times \{w_1, \ldots, w_{I+J}\}$. From the definition of the functionals W_0 and W_1, we know that for each menu $G \in \mathcal{K}(\mathcal{P}(B))$, we have $W_0(G) = W_1(G)$. The uniqueness part of the additive EU representation theorem in DLR now says that the two transformed measures agree on $\mathfrak{U}_B$ and, hence, $\mu_0 = \mu_1$, as desired. *Q.E.D.*

It will be useful to denote Y-trivial, additive, separable EU representations of $\succsim$ by the collection $(\mathcal{U}, v, \mu)$ where $v : Y \to \mathbb{R}$ is a vN-M function and μ is a probability measure on (the Borel σ-algebra of) $\mathcal{U}$. It induces the preference functional

$$V(G) = \int_{\mathcal{U}} \max_{p \in G}\big[u(p_m) + v(p_y)\big]\, d\mu(u) \tag{C.3}$$

that represents the preference $\succsim$ over menus. The transformation from the state space with $(r, \alpha) \in \mathfrak{U}_M \times [0,1]$ to the state space with $u \in \mathcal{U}$ is achieved by setting $u := \frac{\alpha}{(1-\alpha)} r$ and by applying the appropriate transformation to the measure, as in Lemma 16.

APPENDIX D: PROOFS FOR SECTION 3

D.1. *Proof of Proposition 2*

Let $W \in C(H \times S)$ and consider the function $\Phi W(f, s)$, given by

$$\Phi W(f, s) := \sum_{s' \in S} \Pi(s, s') \left[\int_{\mathcal{U}} \max_{p \in f(s')} \big[u(p_m) + \delta W(p_h, s')\big]\, d\mu_{s'}(u)\right]$$

for all $s \in S$. It is easy to see that Φ is monotone, that is, $W \le W'$ implies $\Phi W \le \Phi W'$, and satisfies discounting, that is, $\Phi(W + \rho) \le \Phi W + \delta\rho$ when $\rho \ge 0$. If we assume that $\Phi W \in C(H \times S)$ for all $W \in C(H \times S)$, it follows that Φ is a contraction mapping (with modulus δ) and has a unique fixed point, which establishes the proposition. All that remains is to show that Φ is an operator on $C(H \times S)$.

Before completing the proof, we establish a useful inequality, $\int_{\mathcal{U}} \|u\|_2 \, \mathrm{d}\mu_s < \infty$ for each μ_s. To see that the inequality holds, let us define $W: \mathcal{K}_M \to \mathbb{R}$ as $W(a) := \int_{\mathcal{U}} \max_{\alpha \in a} u(\alpha) \, \mathrm{d}\mu(u)$ for $a \in \mathcal{K}_M$. Recall that p_m^* is the uniform lottery over M. Let $a := \{\alpha \in \mathcal{P}(M) : \|\alpha - p_m^*\|_2 \leq \varepsilon$ for some $\varepsilon > 0\}$ so that $a \in \mathcal{K}_M$. Then, for each $u \in \mathcal{U}$, $\max_{\alpha \in a} u(\alpha) = \varepsilon \|u\|_2$. Therefore, $0 \leq W(a) = \varepsilon \int_{\mathcal{U}} \|u\|_2 \, \mathrm{d}\mu(u)$. But $W(\mathcal{P}(M)) = W(M) \leq \sum_m \mu |u_m| < \infty$ because μ is nice, which implies that $0 \leq W(a) < \infty$.

We now show that if $(f_n) \in H^\infty$ is a sequence that converges to $f \in H$, then $\Phi(W)(f_n, s') \to \Phi(W)(f, s')$ whenever $W \in C(H \times S)$, which establishes that $\Phi W \in C(H \times S)$. (Since S is finite, any convergent sequence in S must eventually be constant, which we take to be s'.)

For each $x \in \mathcal{K}$, $u \in \mathcal{U}$, $W \in C(H \times S)$, and $s \in S$, define $\varphi(x, u, s) = \max_{p \in x}[u(p_m) + \delta W(p_h, s)]$. Then $|\varphi(x, u, s)| \leq \max_{p \in x} |u(p_m) + \delta W(p_h, s)| \leq \|u\|_2 \max_{p \in x} |\frac{u(p_m)}{\|u\|_2}| + \max_{p \in x} \delta |W(p_h, s)| \leq \|u\|_2 K_1 + K_{2,s}$ where $K_{2,s} > 0$, $K_1 := \max_{x \in \mathcal{K}} \max_{p \in x} |\frac{u(p_m)}{\|u\|_2}|$, and the bounds follow from the definition of $u \in \mathcal{U}$, the compactness of H, and the continuity of W.

As W is continuous, the function $u(p_m) + \delta W(p_h, s) \in C(M \times H)$ is a continuous, linear functional on $\mathcal{P}(M \times H)$ when the latter is endowed with the topology of weak convergence (which is metrizable). Therefore, by the Maximum Theorem, for each $u \in \mathcal{U}$ and $s \in S$, $\varphi(x, u, s)$ is continuous in x.

Consider the sequence (f_n) that converges to f. By the bounds established above, $|\varphi(f_n(s), u, s)| \leq \|u\|_2 K_1 + K_{2,s}$ and $\|u\|_2 K_1 + K_{2,s}$ is μ_s-integrable since μ_s is nice (by the inequality established above). Moreover,

$$
\begin{aligned}
\lim_{n \to \infty} \Phi W(f_n, s') &= \lim_{n \to \infty} \sum_s \Pi(s', s) \left[\int_{\mathcal{U}} \varphi(f_n(s), u, s) \, \mathrm{d}\mu_s(u) \right] \\
&= \sum_s \Pi(s', s) \left[\int_{\mathcal{U}} \lim_{n \to \infty} \varphi(f_n(s), u, s) \, \mathrm{d}\mu_s(u) \right] \\
&= \sum_s \Pi(s', s) \left[\int_{\mathcal{U}} \varphi(f(s), u, s) \, \mathrm{d}\mu_s(u) \right] \\
&= \Phi W(f, s').
\end{aligned}
$$

As f and (f_n) are arbitrary, we conclude that $\Phi W \in C(H \times S)$ whenever $W \in C(H \times S)$. The equalities above rely on the Dominated Convergence Theorem to interchange the order of limits and integration, and the continuity of $\varphi(\cdot, u, s)$ for each u and s to establish the pointwise limit. This completes the proof.

D.2. *Proof of Theorem 1*

That the representation satisfies all the axioms is straightforward. Consider a preference $\succsim$ that satisfies Axioms 1–9. By Theorem 3, $\succsim$ has a finitely additive, EU representation. Separability (Axiom 5) implies, according to Lemma 19 and Proposition 20, that any such finitely additive EU representation also has a representation based on a regular countably additive measure.

In Section D.2.1, we use S-CSR (Axiom 8) to establish in Proposition 24 that for each $s \in S$, the preference $\succsim_s$ has an H-trivial, separable, additive representation as in (C.3). This amounts to showing that the relevant state space is isomorphic to $\mathcal{U}$. In Section D.2.2, Proposition 25 establishes that the parameter of the representation of $\succsim_s$, namely the measure μ_s on $\mathcal{U}$, is uniquely identified up to a scaling. In Section D.2.3, we invoke the remainder of the axioms and show in Proposition 28 that $\succsim$ has a recursive representation, possibly with state-dependent discount factors. In Section D.2.4, we show that there exists an equivalent representation of $\succsim$ with a state-independent discount factor $\delta \in (0, 1)$. We also show that the collection of measures (μ_s) in this representation is unique up to a common scaling and that the Markov chain on S with transition probabilities Π has a unique stationary distribution π (Propositions 30, 31, and 32). This proves the theorem.

D.2.1. *Separable Representation*

It is easy to see that $\succsim_s$ is continuous, and satisfies Independence, Monotonicity, and Separability, so that by Proposition 20, $\succsim_s$ has a finitely additive separable representation.

PROPOSITION 24: *Suppose $\succsim_s$ has a finitely additive separable representation. If $\succsim_s$ also satisfies State Contingent CSR (Axiom 8), then the representation is H-trivial.*

PROOF: As in Proposition 20, we know that a separable, finitely additive separable representation has the form

$$U_s(x) = \int_{\mathfrak{U}_M \times [0,1] \times \mathfrak{U}_H} \max_{p \in x} \big[\gamma r(p_m) + (1-\gamma) v(p_h)\big] \, \mathrm{d}\mu_s(r, \gamma, v).$$

For an arbitrary consumption prize $m \in M$, let (m, A) be a *rectangular menu*, defined as $(m, A) := \{(m, p_z) : p_z \in A, A \in \mathcal{K}_H\}$. Define a utility function $W_s : \mathcal{K}_H \to \mathbb{R}$ as $W_s(A) = U_s(m, A)$. It follows from the separability of the representation that the choice of $m \in M$ only affects W_s up to a constant. Let $\mathrm{d}\mu_s^*(v) = \iint_{\mathfrak{U}_M \times [0,1]} \mathrm{d}\mu_s(r, \gamma, v)$ be the marginal of μ on $\mathfrak{U}_H$. Then, $W_s(A) = \int_{\mathfrak{U}_H} \max_{p_h \in A} v(p_h) \, \mathrm{d}\mu_s^*(v) + \text{constant}$.

By Proposition 15 above, CSR (Axiom 8) implies that $W_s(A) = \max_{p_h \in A} v_s(p_h)$, where $v_s \in C(H)$ is given by $v_s(f) = W_s(\{f\})$. But this implies that $\max_{p_h \in A} v_s(p_h) = \int \max_{p_h \in A} v(p_h)\,\mathrm{d}\mu^*(v)$. Therefore, the carrier of μ^* must be a singleton. *Q.E.D.*

It follows from Lemma 22 that μ_s can be extended in a unique way to an additive measure on $\mathcal{U}$. Abusing notation, this transformed measure is also denoted by μ_s. Thus, as mentioned at the end of Section C.3, $\succsim_s$ has a separable, additive, H-trivial EU representation (μ_s, v_s) that can be written as

$$U_s(x) = \int_{\mathcal{U}} \max_{p \in x} \big[u(p_m) + v_s(p_h)\big]\,\mathrm{d}\mu_s(u). \tag{D.1}$$

We end with the observation that μ_s is nice. To see this, note that there exists $p_h^* \in \mathcal{P}(H)$ such that $v_s(p_h^*) = 0$. Now consider the menu (m, p_h^*). It is easy to see that $U_s((m, p_h^*)) = \mu_s u(m) := \int_{\mathcal{U}} u(m)\,\mathrm{d}\mu_s(u)$. But $U_s((m, p_h^*))$ is finite, which implies that $\mu_s u(m)$ is finite for every $m \in M$, which proves that μ_s is nice.

D.2.2. *Identification*

Section D.2.1 shows that the state space that is relevant for an additive, separable, H-trivial representation is $\mathcal{U}$. Our goal in this section is to show that the measure μ_s on $\mathcal{U}$ and the utility function v_s are unique up to a common scaling. We say that v_s is *nondegenerate* if $v_s \neq \mathbf{0}$.

PROPOSITION 25: *Consider two separable, H-trivial, additive EU representations (μ_s, v_s) and (μ'_s, v'_s) of $\succsim_s$, wherein both v_s and v'_s are nondegenerate. There exists $\zeta > 0$ such that the following properties hold*:

(a) $\mu_s(\zeta D) = \mu'_s(D)$ *for all measurable $D \subset \mathcal{U}$*,

(b) $v'_s = \zeta v_s + $ *constant*.

PROOF: By Lemma 16, our result on the change of state space, we know that there exists a measurable bijection $\Psi : \mathcal{U} \to \mathcal{U}$, and integrable functions $(\zeta, \xi) : \mathcal{U} \mapsto (\mathbb{R}_{++}, \mathbb{R})$ such that for each u' in the support of μ'_s, we have $u'(p_m) + v'_s(p_h) = \zeta(u)[(\Psi u)(p_m) + v_s(p_h)] + \xi(u)$.

Consider two lotteries p and q such that $p_m = q_m$ but where $v_s(p_h) \neq v_s(q_h)$. Then

$$\begin{aligned}\big[u'(p_m) + v'_s(p_h)\big] - \big[u'(q_m) + v'_s(q_h)\big] &= v'_s(p_h) - v'_s(q_h) \\ &= \zeta(u)\big[v_s(p_h) - v_s(q_h)\big].\end{aligned}$$

But this implies $\zeta(u) = \frac{v'_s(p_h) - v'_s(q_h)}{v_s(p_h) - v_s(q_h)}$ for all u and, hence, $\zeta(u)$ is constant, which proves (b).

Let U_s' be the functional induced by (v_s', μ_s'). Then

$$\begin{aligned} U_s'(x) &= \int_{\mathcal{U}} \max_{p\in x}\big[u'(p_m) + v_s'(p_h)\big]\,\mathrm{d}\mu_s'(u') \\ &= \int_{\mathcal{U}} \max_{p\in x} \zeta\big[(\Psi u)(p_m) + v_s(p_h)\big]\,\mathrm{d}\tilde{\mu}_s'(u) + \int \xi\mu_s' \\ &= \zeta \int_{\mathcal{U}} \max_{p\in x}\big[\tilde{u}(p_m) + v_s(p_h)\big]\,\mathrm{d}\mu_s''(\tilde{u}) + \text{constant}, \end{aligned}$$

where $\tilde{\mu}_s'$ and μ_s'' each obtain from a change of state space. Proposition 23 says that $\mu_s'' = \mu_s$, which implies that Ψ is the identity mapping. Thus, it must be that $\mu_s'(D) = \mu_s(\zeta D)$ for all measurable $D \subset \mathcal{U}$. These observations prove part (a). *Q.E.D.*

D.2.3. *Recursive Representation*

The preference $\succsim$ on H is continuous and satisfies Independence (Axiom 3). Therefore, there exists an additively separable representation with state-dependent utilities as specified in (D.1), namely

$$W_0(f) = \sum_s \pi(s)\bigg[\int_{\mathcal{U}} \max_{p\in f(s)}\big[u(p_m) + v_s(p_h)\big]\,\mathrm{d}\mu_s(u)\bigg], \tag{D.2}$$

where π is the period 0 beliefs about the states in S. In the representation above, each $v_s : H \to \mathbb{R}$ is continuous. Let $\succsim_s^*$ be the preference relation on H induced by v_s and recall that H is convex.

PROPOSITION 26: *Suppose $\succsim$ has a representation as in (D.2) and satisfies Singleton Indifference to Timing (Axiom 7). Then each $\succsim_s^*$ satisfies Independence.*

PROOF: Let $f, g \in H$ be such that $v_s(f) > v_s(g)$, which implies $U_s(\{p_m^*, f\}) > U_s(\{p_m^*, g\})$. Let $f' \in H$ and fix $\lambda \in (0, 1]$. By the linearity of U_s, we have $U_s(\lambda\{p_m^*, f\} + (1-\lambda)\{p_m^*, f'\}) > U_s(\lambda\{p_m^*, g\} + (1-\lambda)\{p_m^*, f'\})$. But $\succsim_s$ satisfies Singleton Indifference to Timing (Axiom 7), which implies that $U_s(\lambda\{p_m^*, f\} + (1-\lambda)\{p_m^*, f'\}) = U_s(\{p_m^*, \lambda f + (1-\lambda)f'\})$, so that $U_s(\{p_m^*, \lambda f + (1-\lambda)f'\}) > U_s(\{p_m^*, \lambda g + (1-\lambda)f'\})$, which in turn implies that $v_s(\lambda f + (1-\lambda)f') > v_s(\lambda g + (1-\lambda)f')$. We have just established that $\succsim_s^*$ satisfies Independence (Axiom 3). *Q.E.D.*

PROPOSITION 27: *Suppose $\succsim$ has a representation as in (D.2). If $\succsim$ satisfies Statewise Nontriviality (Axiom 1), History Independence (Axiom 9), and Aggregate Stationarity (Axiom 6), then each $\succsim_s^*$ also satisfies Monotonicity and Statewise Nontriviality (Axiom 1).*

PROOF: By Proposition 26, $\succsim_s^*$ satisfies Independence (Axiom 3) and, hence, v_s is linear. It follows that there exist utility functions $w_{s,s'}: \mathcal{K} \to \mathbb{R}$ for all $s' \in S$ such that $v_s(f) = \sum_{s'} w_{s,s'}(f(s'))$. Notice that by History Independence (Axiom 9), we may assume that $w_{s,s'} = \pi_s(s')w_{s'}$ for all $s, s' \in S$, where $\pi_s(s') > 0$.

Notice that $f_s^x \succsim f_s^y$ if and only if $U_s(x) \geq U_s(y)$. By Aggregate Stationarity (Axiom 6), this is equivalent to requiring that $\{(m, f_s^x)\} \succsim \{(m, f_s^y)\}$, which holds if and only if $\sum_s' \pi(s')v_{s'}(f_s^x) \geq \sum_s' \pi(s')v_{s'}(f_s^y)$. By the observations above, this is equivalent to $\sum_{s'} \pi(s')\pi_{s'}(s)w_s(x) \geq \sum_{s'} \pi(s')\pi_{s'}(s)w_s(y)$. Let $W_s'(x) := \sum_{s'} \pi(s')\pi_{s'}(s)w_s(x)$. Then W_s' is linear on $\mathcal{K}$ and $W_s'(x) \geq W_s'(y)$ if and only if $U_s(x) \geq U_s(y)$.

Thus, $\succsim_s^*$ satisfies Monotonicity (Axiom 4) if and only if $w_s(x \cup y) \geq w_s(x)$ for all $x, y \in \mathcal{K}$. But w_s is monotone with respect to set inclusion if and only if U_s is, which proves that $f_{s'}^{x\cup y} \succsim_s^* f_{s'}^x$ because $x \cup y \succsim_s x$.

Finally, notice that $W_{\tilde{s}}' = \sum_s \pi(s)\pi_s(\tilde{s})w_{\tilde{s}}$. But $W_{\tilde{s}}'$ is nontrivial because $\succsim_{\tilde{s}}$ is nontrivial. Therefore, it must be that $w_{\tilde{s}}$ is nontrivial, from which it follows that $\succsim_{\tilde{s}}^*$ satisfies Statewise Nontriviality (Axiom 1). *Q.E.D.*

We now establish the existence of a *recursive* representation.

PROPOSITION 28: *Let $\succsim$ have a representation as in* (D.2), *and suppose $\succsim$ satisfies Statewise Nontriviality* (*Axiom* 1), *History Independence* (*Axiom* 9), *and Singleton Indifference to Timing* (*Axiom* 7). *Then there is a value function*

$$V(f, s') = \sum_s \Pi(s', s)\left[\int_{\mathcal{U}} \max_{p\in f(s)}\left[u(p_m) + \delta_s V(p_h, s)\right] \mathrm{d}\mu_s(u)\right], \tag{D.3}$$

where Π governs transition probabilities for a Markov process on S and $\Pi(s', s) > 0$ for all $s', s \in S$, such that

$$V_0(f) = \sum_s \pi(s)\left[\int_{\mathcal{U}} \max_{p\in f(s)}\left[u(p_m) + \delta_s V(p_h, s)\right] \mathrm{d}\mu_s(u)\right] \tag{D.4}$$

represents $\succsim$.

PROOF: Fix $s' \in S$ and consider the act $f_{s'}^x$. Then

$$W_0(f_{s'}^x) = \pi(s') \int_{\mathcal{U}} \max_{p\in x}\left[u(p_m) + v_{s'}(p_h)\right] \mathrm{d}\mu_{s'}(u) + (\cdot).$$

As noted above, $v_s(\cdot)$ induces a preference $\succsim_s^*$ on H such that (i) $\succsim_s^*$ is continuous, (ii) $\succsim_s^*$ satisfies Independence (Proposition 26), and (iii) $\succsim_s^*$ satisfies Monotonicity (Proposition 27). Then, by the Mixture Space Theorem and by

Theorem 3, v_s has a (state separable) finitely additive representation, which can be written as

$$v_s\big(f^x_{s'}\big) = \delta_s \pi_s\big(s'\big)\int_{\mathfrak{U}_{M\times H}} \max_{p\in x} \mathfrak{u}(p)\,\mathrm{d}\tilde{\mu}^s_{s'}(\mathfrak{u}) + \delta_s \sum_{t\neq s'} \pi_s(t)\int_{\mathfrak{U}_{M\times H}} \max_{p\in f(t)} \mathfrak{u}(p)\,\mathrm{d}\tilde{\mu}^s_t(\mathfrak{u}),$$

where δ_s is a scaling factor chosen so that (i) $\pi_s \in \mathcal{P}(S)$ is strictly positive and (ii) $\tilde{\mu}^s_t$ is a probability charge for all $t \in S$. Such a choice can be made as follows: If $\tilde{\mu}^s_t(\mathfrak{U}_{M\times H}) > 1$, define $\pi'_s(t) := \tilde{\mu}^s_t(\mathfrak{U}_{M\times H})$, so that we can take $\tilde{\mu}^s_t$ to be a probability measure. Because $\pi'_s(t) > 0$ for all $t \in S$, it follows that we may let $\delta_s = \sum_t \pi'_s(t)$, so that the scaling factor $\pi'_s(t)$ can be replaced by $\pi_s(t) := \pi'_s(t)/\delta_s$.

Singleton Indifference to Timing (Axiom 7) implies that

$$\pi(s')\int_{\mathcal{U}} \max_{p\in x}\big[u(p_m) + v_{s'}(p_h)\big]\,\mathrm{d}\mu_{s'}(u) \quad \text{and}$$

$$\delta_s \pi_s(s')\int_{\mathfrak{U}_{M\times H}} \max_{p\in x} \mathfrak{u}(p)\,\mathrm{d}\tilde{\mu}^s_{s'}(\mathfrak{u})$$

are (positive) affine transformations of each other. By Statewise Nontriviality (Axiom 1) every state is nonnull under π, and by History Independence (Axiom 9), every state is also nonnull under π_s, that is, $\pi_s(s') > 0$ (this was established in Proposition 27). Hence the measures $\{\pi, \pi_s : s \in S\}$ have full support.

Recall that $\mu_{s'}$ defined on $\mathcal{U}$ is the marginal of a measure on $\mathcal{U} \times \mathfrak{U}_H$, where the marginal on $\mathfrak{U}_H$ has a singleton carrier. The transformations of the state space used to obtain a separable representation are an instance of those considered in Lemma 16, so that we may regard $\mu_{s'}$ as a measure on $\mathfrak{U}_{M\times H}$.

Since both $\tilde{\mu}^s_{s'}$ and $\mu_{s'}$ are probability charges, we must have $\tilde{\mu}^s_{s'} = \mu_{s'}$ for all $s \in S$. Therefore, each v_s can be written as

$$v_s(f) = \delta_s \sum_{s'} \pi_s\big(s'\big)\int_{\mathcal{U}} \max_{p\in f(s')}\big[u(p_m) + v_{s'}(p_h)\big]\,\mathrm{d}\mu_{s'}(u).$$

Define the Markov transition probabilities on S to be $\Pi(s, s') := \pi_s(s')$ and let $V(f, s) := v_s(f)/\delta_s$ so that

$$V(f, s) = \sum_{s'} \Pi\big(s, s'\big)\int_{\mathcal{U}} \max_{p\in f(s')}\big[u(p_m) + \delta_{s'} V\big(p_h, s'\big)\big]\,\mathrm{d}\mu_{s'}(u).$$

Finally, define

$$V_0(f) := W_0(f) = \sum_s \pi(s) \int_{\mathcal{U}} \max_{p \in f(s)} \big[u(p_m) + \delta_s V(p_h, \pi_s)\big] \, \mathrm{d}\mu_s(u)$$

to attain the desired recursive representation. *Q.E.D.*

It follows from the recursive representation and the fact that $\succsim$ satisfies Statewise Nontriviality (Axiom 1) that at least one μ_s must be nontrivial. To see this, suppose every μ_s is trivial, that is, $\mu_s(\{\mathbf{0}\})$ for all $s \in S$. Then the recursive representation implies every $V(\cdot, s)$ must be identically 0, which violates Statewise Nontriviality (Axiom 1).

D.2.4. *Unique Representation With a State-Independent Discount Factor*

Proposition 28 establishes the existence of a recursive value function as in (D.3), which can be described by the parameters $((\mu_s, \pi_s, \delta_s)_{s \in S})$. We now show that there exists a unique equivalent representation with a state-*independent* discount factor. Let $\xi \in \mathbb{R}^S_{++}$ and let $\langle \pi, \xi \rangle := \sum_s \pi(s)\xi(s)$ for any $\pi \in \mathcal{P}(S)$.

PROPOSITION 29: *Let $((\mu_s, \pi_s, \delta_s)_{s \in S})$ and $((\hat{\mu}_s, \hat{\pi}_s, \hat{\delta}_s)_{s \in S})$ be two recursive representations of $\succsim$ as in (D.3). Then there exists $\xi \in \mathbb{R}^S_{++}$ such that the parameters are related by the transformations $\hat{\pi}(s) := \pi(s)\xi(s)/\langle \pi, \xi \rangle$, $\hat{\pi}_s(s') := \pi_s(s')\xi(s')/\langle \pi_s, \xi \rangle$, $\hat{\mu}(D) := \mu(D/\xi(s))$, $\hat{\delta}_s := \delta_s \langle \pi_s, \xi \rangle / \xi(s)$, and $\hat{V}(\cdot, \hat{\pi}_s) := V(\cdot, \pi_s)/\langle \pi_s, \xi \rangle$. Moreover, if $((\mu_s, \pi_s, \delta_s)_{s \in S})$ is a representation of $\succsim$, then the transformed set of parameters $((\hat{\mu}_s, \hat{\pi}_s, \hat{\delta}_s)_{s \in S})$ also represents $\succsim$.*

PROOF: The last part of the proposition is immediate. By our identification result, Proposition 25, there exists $\xi \in \mathbb{R}^S_{++}$ such that (i) $\hat{\mu}_s(D) = \mu_s(D/\xi(s))$ for all $s \in S$ and for all Borel measurable $D \subset \mathcal{U}$, and (ii) $\hat{\delta}_s \hat{V}(\cdot, s) := \delta_s V(\cdot, s)/\xi(s)$. By construction, $\hat{V}(\cdot, s) = V(\cdot, s)/\langle \pi_s, \xi \rangle$ and, hence, we must have $\hat{\delta}_s := \delta_s \langle \pi_s, \xi \rangle / \xi(s)$. Finally, the identification from the Mixture Space Theorem implies that $\hat{\pi}_s(s') := \pi_s(s')\xi(s')/\langle \pi_s, \xi \rangle$. *Q.E.D.*

PROPOSITION 30: *There exists $\xi \in \mathbb{R}^S_{++}$ such that $\hat{\delta}_s$ is independent of $s \in S$. Moreover, ξ is unique up to scaling, so that $\hat{\delta}$ is unique, and the corresponding measures $(\hat{\mu}_s)_{s \in S}$ are unique up to scaling.*

PROOF: A representation with a constant discount factor will obtain immediately if we can establish that there exists a vector $\xi \gg \mathbf{0}$ and a number $\hat{\delta} > 0$ such that $\hat{\delta}\xi(s) = \delta_s \langle \pi_s, \xi \rangle$ for all $s \in S$.

For $S = \{1, \dots, n\}$, consider the stochastic matrix Π, whose row s is π_s. Define the $n \times n$ diagonal matrix Δ as $\Delta(i, j) = \delta_i$ if $j = i$ and 0 otherwise. In matrix notation, our problem amounts to finding a $\xi \gg \mathbf{0}$ and $\hat{\delta} > 0$ such that $\hat{\delta}\xi = \xi\Delta\Pi$. This amounts to showing that (i) ξ is a (left) eigenvector of the matrix $\Delta\Pi$ and (ii) $\hat{\delta}$ is the corresponding eigenvalue.

We say that a matrix is *positive* if each of its entries is strictly positive. By Proposition 28 above, the matrix Π is positive, so the matrix $\Delta\Pi$ is also positive. Thus, by the Perron theorem below, such ξ and $\hat{\delta}$ exist, ξ is unique up to scaling, and $\hat{\delta}$ is unique.

By Proposition 29, these are the only transformations that we need consider, which establishes that the measures $(\hat{\mu}_s)_{s\in S}$ are unique up to a common scaling. *Q.E.D.*

The Perron Theorem is standard and can be found, for instance, as Theorem 1 in Chapter 16 of Lax (2007). For completeness, we state the relevant part of the theorem.

THEOREM 4—Perron: *Every positive matrix* A *has a* dominant eigenvalue *denoted by* $\hat{\delta}$, *which has the following properties*:

(a) $\hat{\delta} > 0$ *and the associated eigenvector* $\xi \gg \mathbf{0}$.

(b) $\hat{\delta}$ *is a simple eigenvalue and, hence, it has algebraic and geometric multiplicity* 1.

(c) A *has no other eigenvector with nonnegative entries*.

Thus, we have established the existence of a recursive value function

$$V(f, s') = \sum_s \Pi(s', s)\left[\int_{\mathcal{U}} \max_{p\in f(s)}\left[u(p_m) + \delta V(p_h, s)\right] \mathrm{d}\mu_s(u)\right] \tag{D.5}$$

such that

$$V_0(f) = \sum_s \pi(s)\left[\int_{\mathcal{U}} \max_{p\in f(s)}\left[u(p_m) + \delta V(p_h, s)\right] \mathrm{d}\mu_s(u)\right] \tag{D.6}$$

represents $\succsim$.

PROPOSITION 31: *Let* $\succsim$ *have a recursive representation as in* (D.6). *If* $\succsim$ *satisfies Aggregate Stationarity* (*Axiom* 6), *then* π *is the unique stationary distribution of the Markov process with transition matrix* Π.

PROOF: Recall that $V_0(\{(m, f)\}) = \sum_s \pi(s)[\int u(m) + \delta V(f, s)]\,\mathrm{d}\mu_s(u)$. Then, letting $\kappa := \sum_s \pi(s)\mu_s u(m)$, we see that

$$\begin{aligned} &V_0(\{(m, f)\}) - \kappa \\ &\quad = \delta \sum_s \pi(s)\Big[\sum_{s'} \Pi(s, s')\Big[\int_{\mathcal{U}} \max_{p\in f(s')}\big[u(p_m) + \delta V(p_h, s')\big]\,\mathrm{d}\mu_{s'}(u)\Big]\Big] \\ &\quad = \delta \sum_{s'}\Big[\sum_s \pi(s)\Pi(s, s')\Big]\Big[\int_{\mathcal{U}} \max_{p\in f_s^x}\big[u(p_m) + \delta V(p_h, s')\big]\,\mathrm{d}\mu_{s'}(u)\Big]. \end{aligned}$$

By Aggregate Stationarity (Axiom 6), $V_0(\{(m, \cdot)\})$ and $V_0(\cdot)$ represent the same preference. From Propositions 26 and 28, we see that each $V(\cdot, s)$ is linear on H. Therefore, $V_0(\{(m, \cdot)\})$ is linear over H. By the Mixture Space Theorem, $V_0(\{(m, \cdot)\})$ and $V_0(\cdot)$ must be affine transformations of each other. Therefore, $\pi(s') = \sum_s \pi(s)\Pi(s, s')$ for all $s' \in S$. In other words, π is a stationary distribution of Π. Since Π is positive, the stationary distribution is unique. *Q.E.D.*

PROPOSITION 32: *If $\succsim$ has a recursive representation of the form in (D.6), then* $\delta \in (0, 1)$.

PROOF: It follows immediately from the nontriviality of $\succsim$ and from Aggregate Stationarity (Axiom 6) that $\delta > 0$. We now show that $\delta < 1$.

As π is the unique stationary distribution of Π, we have for any $f \in H$, $V_0(f) = \sum_s \pi(s)V(f, s)$. Now let $a \in \mathcal{K}_M$ be the closed ε-ball around the p_m^*, the uniform lottery over M. It is clear that $0 < \int_{\mathcal{U}} \max_{p\in a} u(p_m)\,\mathrm{d}\mu_s(u)$. But we also have $\int_{\mathcal{U}} \max_{p\in a} u(p_m)\,\mathrm{d}\mu_s(u) \le \int_{\mathcal{U}} \max_{p\in \mathcal{P}(M)} u(p_m)\,\mathrm{d}\mu_s(u) = \int_{\mathcal{U}} \max_{m\in M} u(m)\,\mathrm{d}\mu_s(u) \le \sum_{m\in M} |\mu_s u(m)| < \infty$, where the last inequality holds because μ_s is nice. By the recursive construction of H, it follows that there exists a unique act f^* that gives the menu a in each period and in every state, that is, $f^* \simeq (a, f^*)$. Letting $\eta := \sum_s \pi(s)[\int_{\mathcal{U}} \max_{p_m\in a} u(p_m)\,\mathrm{d}\mu_s(u)]$, we see that $V_0(f^*) = \eta + \delta V_0(f^*) = \eta \sum_{\tau\ge 0} \delta^\tau$. Since $\eta > 0$ and because $V_0(f^*)$ is finite, we conclude that $\delta < 1$, as required. *Q.E.D.*

APPENDIX E: PROOFS FOR SECTION 4

PROOF OF COROLLARY 4: Suppose first that $\succsim$ has a DPF representation $((\mu_s)_{s\in S}, \Pi, \delta)$ and that each $\succsim_s|_{\mathcal{K}(\mathcal{P}(M\times\{f^*\}))}$ is strategically rational. By Proposition 15, strategic rationality implies that the support of each μ_s lies in $\{\lambda u_s : \lambda \ge 0\}$ for some $u_s \in \mathcal{U}$.

For the converse, suppose the support of μ_s lies in $\{\lambda u_s : \lambda \ge 0\}$ for some $u_s \in \mathcal{U}$. It is easy to see that for any $a \in \mathcal{K}(\mathcal{P}(M))$, $\int_{\mathbb{R}_+} \max_{\alpha\in a} \lambda u_s(\alpha)\,\mathrm{d}\mu(\lambda) =$

$\max_{\alpha\in a}\mu_s u(\alpha)$. Clearly, the functional $a \mapsto \max_{\alpha\in a}\mu_s u(\alpha)$ represents $\succsim_s |_{\mathcal{K}(\mathcal{P}(M\times\{f^*\}))}$, from which it follows that the latter is strategically rational. *Q.E.D.*

PROOF OF COROLLARY 5: We define a sequence of extensions. Let $H_0 := \mathcal{H}(\mathcal{K}(\mathcal{P}(M \times L)))$ and inductively define $H_n := \mathcal{H}(\mathcal{K}(\mathcal{P}(M \times H_{n-1})))$. It is easy to see that $H_{n-1} \subset H_n$ and $H_n \uparrow H$.

For each $x \in \mathcal{K}(\mathcal{P}(M \times L))$, let $\varphi_0(x) := \{p : p \gtrsim q\ \forall q' \in x\}$. Similarly, for each $f \in H_0$, let $\psi_0(f)(s) := q$, where $q \in \varphi_0(f(s))$, and notice that $\psi_0(f) \in L$. Now define $\succsim^0$ on H_0 as $f \succsim^0 g$ if and only if $\psi_0(f) \gtrsim \psi_0(g)$.

Now suppose $\succsim_{n-1}$ has been defined on H_{n-1}. For each $x \in \mathcal{K}(\mathcal{P}(M \times H_{n-1}))$, let $\varphi_{n-1}(x) := \{p : p \succsim_{n-1} q\ \forall q' \in x\}$. Similarly, for each $f \in H_{n-1}$, let $\psi_{n-1}(f)(s) := q$, where $q \in \varphi_{n-1}(f(s))$, and notice that $\psi_{n-1}(f) \in H_{n-1}$. Now define $\succsim^n$ on H_n as $f \succsim^n g$ if and only if $\psi_{n-1}(f) \succsim^{n-1} \psi_{n-1}(g)$.

Let $\succsim = \lim_{n\to\infty} \succsim^n$. Notice that this limit is unique because the extension at each stage is unique, and because each preference is continuous and H is compact. Notice also that $\succsim$ is a well defined preference relation on H and by construction satisfies Axioms 1–9, and each $\succsim_s$ is strategically rational. This implies that $\succsim$ has a DPF representation $((\mu_s)_{s\in S}, \Pi, \delta)$, where, by the strategic rationality of $\succsim_s$, each μ_s is a Dirac measure.

Now suppose, contrary to the hypothesis, that there exist two recursive Anscombe–Aumann representations of $\gtrsim$ on L, namely $((u_s), \Pi, \delta)$ and $((u_s^*), \Pi^*, \delta^*)$. Let μ_s be the Dirac measure on $\mathcal{U}$ concentrated on u_s and let μ_s^* be the Dirac measure concentrated on u_s^* for each s. Then $((\mu_s), \Pi, \delta)$ and $((\mu_s^*), \Pi^*, \delta^*)$ are two DPF representations that are identical on L.

It is easy to see that both $((u_s), \Pi, \delta)$ and $((u_s^*), \Pi^*, \delta^*)$ represent each $\succsim^n$ on H_n for all $n \geq 0$. Therefore, both represent $\succsim$ on H. But by the uniqueness of the DPF representation, this implies $\Pi = \Pi^*$, $\delta = \delta^*$, and $(\mu_s)_{s\in S}$ and $(\mu_s^*)_{s\in S}$ are identical up to a common scaling. This establishes the uniqueness of the recursive Anscombe–Aumann representation, as required. *Q.E.D.*

PROOF OF COROLLARY 6: Let us suppose that μ_s is independent of s. Then, clearly, $V(\ell, s)$ is independent of s for all $\ell \in L_0$. But this implies that for all $x \in X$, $\int \max_{p\in x}[u(p_m) + V(p_\ell, s)]\,\mathrm{d}\mu_s(u)$ is independent of s. This implies that $\succsim_s |_X$ is independent of s, as claimed.

For the converse, let us assume that $\succsim_s |_X$ is independent of s. Then, for any $s, s' \in S$, it is the case that for all $m \in M$ and $\ell, \ell' \in L_0$, $(m, \ell) \succsim_s (m, \ell')$ if and only if $(m, \ell) \succsim_{s'} (m, \ell')$. This implies $V(\ell, s) \geq V(\ell', s)$ if and only if $V(\ell, s') \geq V(\ell', s')$. Therefore, we must have $V_s = \lambda_{s,s'} V_{s'}$ for some $\lambda_{s,s'} > 0$. Let $\ell \in L_0$ deliver lottery p in the first period and p_m^* in every subsequent period. Then $V(\ell, s) = \sum_s \pi(s)\bar{u}_s(p) = V(\ell, s') = \lambda_{s,s'} V(\ell, s)$, which implies $\lambda_{s,s'} = 1$. Therefore, V_s is independent of s on L_0.

As $\succsim_s |_X$ is independent of s, we have $U_s|_{L_0} \propto U_{s'}|_{L_0}$. But we have just established that $V_s|_{L_0}$ is independent of s, which implies that $U_s|_{L_0}$ is also independent of s. Therefore, μ_s must be independent of s, which completes the proof. *Q.E.D.*

APPENDIX F: PROOFS FOR SECTION 5

Lemma 33 is not part of the proof of a result in the text, but it is referred to in Section 5.1.

LEMMA 33: *Let $((u_s)_{s\in S}, \Pi, \delta)$ and $((u'_s)_{s\in S}, \Pi', \delta')$ be DPF representations of $\succsim$, and suppose $\bar{u}'_s = \lambda_s \bar{u}_s$. Then $\lambda_s = \lambda_{s'}$ for all $s, s' \in S$.*

PROOF: Let us suppose λ_s is not independent of s. Let $s^*, s_* \in S$ be such that $\lambda_{s_*} \leq \lambda_s \leq \lambda_{s^*}$ for all $s \in S$ with $\lambda_{s_*} < \lambda_{s^*}$. Given that DM is in state s^* today, let α and β be lotteries over M such that DM is indifferent between (i) receiving α today and p^*_m for sure in all future periods, and (ii) receiving β in state s tomorrow and p^*_m in all other states and periods. Since p^*_m yields zero consumption utility and since the representations are separable and recursive, we have $\bar{u}_{s^*}(\alpha) = \delta\Pi(s^*, s)\bar{u}_s(\beta)$, and $\bar{u}'_{s^*}(\alpha) = \delta'\Pi'(s^*, s)\bar{u}'_s(\beta)$ (where we have used the fact that $\Pi(s, s') > 0$ for all $s, s' \in S$ and similarly for Π'). With $\bar{u}'_s = \lambda_s \bar{u}_s$ and $\lambda_{s^*} \geq \lambda_s$ for all s, we obtain $\delta\Pi(s^*, s) \leq \delta'\Pi'(s^*, s)$. As this holds for all $s \in S$, and because $\Pi(s^*, \cdot)$ and $\Pi'(s^*, \cdot)$ are probability measures, we note that $\delta = \delta\sum_s \Pi(s^*, s) < \delta'\sum_s \Pi'(s^*, s_*) = \delta'$, where the inequality is strict because $\delta\Pi(s^*, s_*) < \delta'\Pi'(s^*, s_*)$.

Similarly, for DM in state s_* with $\lambda_{s_*} \leq \lambda_s$ for all s, we find that $\delta > \delta'$. Therefore, it must be that λ_s is independent of $s \in S$, such that $\bar{u}'_s = \lambda\bar{u}_s$ for all $s \in S$ and for some $\lambda \in \mathbb{R}_+$.

This implies that for all $s, s' \in S$, $\delta\Pi(s, s') = \delta'\Pi'(s, s')$. Because $\Pi(s, \cdot)$ and $\Pi'(s, \cdot)$ are probability measures, it follows that $\delta = \delta'$ and that $\Pi(s, \cdot) = \Pi'(s, \cdot)$ for all $s \in S$, establishing the uniqueness claimed. *Q.E.D.*

We now present some proofs concerning the behavioral comparison "greater preference for flexibility" in Section 5.3. We begin with some preliminary lemmas.

LEMMA 34: *Suppose $\succsim^*$ has a greater preference for flexibility than $\succsim$, and both $\succsim$ and $\succsim^*$ satisfy Independence (Axiom 3). Then, for all $\ell, \ell' \in L$, $\ell \succsim^* \ell'$ if and only if $\ell \succsim \ell'$.*

PROOF: By hypothesis, $\ell \succsim \ell'$ implies $\ell \succsim^* \ell'$ or, equivalently, $\ell' \succ^* \ell$ implies $\ell' \succ \ell$. Therefore, it suffices to show that $\ell' \sim^* \ell$ implies $\ell' \sim \ell$. So let us suppose $\ell' \sim^* \ell$ and, without loss of generality, assume that $\ell' \succ \ell$. Let $\ell^\dagger \succ^* \ell' \sim^* \ell$ (by Independence (Axiom 3), it suffices to consider ℓ and ℓ' for which such an $\ell^\dagger$ exists) so that for some sufficiently small $\lambda \in (0, 1)$, we have $\lambda\ell^\dagger + (1-\lambda)\ell \succ^* \ell'$ but $\ell' \succ \lambda\ell^\dagger + (1-\lambda)\ell$, which contradicts the hypothesis, thereby completing the proof. *Q.E.D.*

LEMMA 35: *Suppose* $\succsim^*$ *has a greater preference for flexibility than* $\succsim$, *and suppose* $\succsim$ *and* $\succsim^*$ *have canonical DPF representations* $((\mu_s)_{s\in S}, \Pi, \delta)$ *and* $((\mu_s^*)_{s\in S}, \Pi^*, \delta^*)$, *respectively*. *Then* $\delta = \delta^*$, $\Pi = \Pi^*$, $\mu_s u = \mu_s^* u$ *for all* s, *and* $V(\ell, s) = V^*(\ell, s)$ *for all* $\ell \in L$ *and* $s \in S$. *Moreover*, $V^*(f, s) \geq V(f, s)$ *for all* $f \in H$.

PROOF: By Lemma 34, we know that $\ell \succsim^* \ell'$ if and only if $\ell \succsim \ell'$. Thus, $\succsim$ and $\succsim^*$ represent the same preference on the restricted domain L. Let V_L and V_L^* denote the value functions for the respective canonical DPF representations, restricted to L. As V_L and V_L^* represent the same preference, they are affine transformations of each other (in each state s). By Corollary 5, $\mu_s u = \lambda \mu_s^* u$ for some $\lambda > 0$, where λ is independent of s. But both DPF representations are canonical, which means that, in fact, $\lambda = 1$.

Let ℓ^* denote the S-IHCS that gives the uniform lottery in every state in each period, so $V_L(\ell^*, s) = V_L^*(\ell^*, s) = 0$ for all s (and hence $V_0(\ell^*) = 0 = V_0^*(\ell^*)$). For any probability measure μ_s on $\mathcal{U}$ with $\|\mu_s u\|_2 \neq 0$, there exist $m, m' \in M$ such that $\mu u(m) > 0 > \mu u(m')$. This implies that there exist $\ell, \ell' \in L$ such that $\ell \succ \ell^* \succ \ell'$, where ℓ^* is as above. Consequently, for any $f \in H$, there exist $\lambda \in (0, 1)$ and $\ell^\dagger \in L$ such that $\lambda f + (1 - \lambda)\ell^* \sim \ell^\dagger$, which means that $\lambda f + (1 - \lambda)\ell^* \succsim^* \ell^\dagger$ because $\succsim^*$ has a greater preference for flexibility than $\succsim$. This implies $V^*(\lambda f + (1 - \lambda)\ell^*, s) \geq V^*(\ell^\dagger, s) = V(\ell^\dagger, s) = V(\lambda f + (1 - \lambda)\ell^*, s)$, from which it follows that $V^*(f, s) \geq V(f, s)$ for all $f \in H$ and $s \in S$, as claimed. *Q.E.D.*

It follows easily from Lemma 34 that for each s, $U_s^* \geq U_s$, where U_s represents $\succsim_s$ and U_s^* represents $\succsim_s^*$.

Let V be the value function that corresponds to the canonical DPF representation $((\mu_s), \Pi, \delta)$. Notice that since there is no preference for flexibility with respect to continuation problems in any state s, for the purpose of assigning utilities, we may restrict attention to the domain $\mathcal{H}(\mathcal{K}(\mathcal{P}(M \times \{h_\circ, h^\circ\})))$, where $V(h_\circ) \leq V(h) \leq V(h^\circ)$ for all $h \in H$, and $V(h_\circ, s) < V(h^\circ, s)$ for all s.

LEMMA 36: *Let* $\varphi^s : \mathcal{U} \to \mathbb{R}$ *be convex, and let* $D \subset \mathcal{U}$ *be compact and convex such that* $\varphi^s|_D$ *is Lipschitz. Then there exists a compact* $x \subset \operatorname{aff} \mathcal{P}(M \times \{h_\circ, h^\circ\})$ *such that* $\max_{p\in x}[u(p_m) + \delta V(p_h, s)] \leq \varphi^s(u)$ *with equality for all* $u \in D$.

PROOF: Let $u^* \in D$ and let $\sigma_{u^*}(u) : \mathcal{U} \to \mathbb{R}$ be an affine function such that $\sigma_{u^*} \leq \varphi^s$ with equality at u^*. (The function σ_{u^*} represents the hyperplane that supports φ^s at u^*.) Then there exist $d \in \mathbb{R}^M$ and $d' \in \mathbb{R}$ such that $\sigma_{u^*}(u) = \langle d, u\rangle + d'$. We now construct the corresponding menu.

Notice that $\sigma_{u^*}(u) = \sum_i d_i u_i + d' = \sum_{i=1}^{M-1} u_i(d_i - d_M) + d'$, where we have used the fact that $\sum_i u_i = 0$ for all $u \in \mathcal{U}$. Let $\alpha \in \operatorname{aff} \mathcal{P}(M)$ be such that $\alpha_i = d_i - d_M$ for $i = 1, \ldots, M - 1$ and $\alpha_M = 1 - \sum_{i=1}^{M-1} \alpha_i$. Also, let $q_h \in \operatorname{aff} \mathcal{P}(\{h_\circ, h^\circ\})$ be such that $d' = \delta V(q_h, s)$ (where V has been extended to aff $\mathcal{P}(\{h_\circ, h^\circ\})$ by linearity and, hence, uniquely). Now consider $q(u^*) = (\alpha, q_h)$, the signed measure on $M \times \{h_\circ, h^\circ\}$ (with marginals α and q_h). Then $\sigma_{u^*}(u) = u(\alpha) + \delta V(q_h, s)$.

Recall that φ^s restricted to D is Lipschitz. This implies that the set $x := \overline{\text{conv}}\{q(u) : u \in D\}$ is compact. (Intuitively, the set of "slopes" and intercepts of the supporting hyperplanes of $\varphi^s|_D$ is compact.) It is clear that $x \subset \text{aff}\, \mathcal{P}(M \times \{h_\circ, h^\circ\})$. For each $q \in x$, $u(q_m) + \delta V(q_h, s)$ is an affine function of u. Therefore, $\max_{q \in x}[u(q_m) + \delta V(q_h, s)]$ is a convex function of u. By construction, this function is always dominated by φ^s and is equal to φ^s on D, which completes the proof. *Q.E.D.*

The proof of Theorem 2 relies on the following theorem, which characterizes dilations.

THEOREM 5: *Let μ and μ^* be probability measures on $\mathcal{U}$. Then the following statements are equivalent*:

(a) *μ^* is a dilation of μ.*

(b) *We have $\mu^* \varphi \geq \mu \varphi$ for every continuous convex function $\varphi : \mathcal{U} \to \mathbb{R}$ that is $\mu + \mu^*$ integrable.*

Theorem 5 is Theorem 7.2.17 in Torgersen (1991). It is proved by Blackwell (1953) for the case where the supports of μ and μ^* are bounded.

PROOF OF THEOREM 2: "Only if": Suppose $\succsim^*$ has greater preference for flexibility than $\succsim$. By Theorem 5, it suffices to show that for every continuous convex function $\varphi^s : \mathcal{U} \to \mathbb{R}$ that is $\mu_s + \mu_s^*$ integrable, we have $\mu_s^* \varphi^s \geq \mu_s \varphi^s$.

Fix such a function φ^s. We show that there exist continuous, convex functions $\varphi_n^s : \mathcal{U} \to \mathbb{R}$ that are $\mu_s + \mu_s^*$ integrable, such that (i) $\varphi_n^s \leq \varphi_{n+1}^s$, (ii) $\varphi_n^s \uparrow \varphi^s$, and (iii) $\mu_s \varphi_n^s \leq \mu_s^* \varphi_n^s$. Since φ_1^s is $(\mu_s + \mu_s^*)$ integrable, it follows that $\varphi^s - \varphi_1^s$ is integrable. Therefore, $0 \leq \varphi^s - \varphi_n^s \leq \varphi^s - \varphi_1^s$, so the dominated convergence theorem and (iii) above imply $\mu_s \varphi^s \leq \mu_s^* \varphi^s$.

Let (D_n) be an increasing sequence of compact subsets of $\mathcal{U}$ such that $\bigcup_n D_n$ covers the effective domain of φ^s. By standard arguments (because the effective domain has a relative interior in our finite dimensional setting), we may assume that for all n, $\varphi^s|_{D_n}$ is Lipschitz. Then, by Lemma 36, there exists $x_n \subset \text{aff}\, \mathcal{P}(M \times \{h_\circ, h^\circ\})$ so that $\varphi_n^s(u) := \max_{q \in x_n}[u(q_m) + \delta V(q_h, s)]$ satisfies $\varphi_n^s \leq \varphi^s$ and $\varphi_n^s|_{D_n} = \varphi^s|_{D_n}$. It follows immediately from the construction (in Lemma 36) that $\varphi_n^s \leq \varphi_{n+1}^s$ (since $x_n \subset x_{n+1}$). All that remains is to show that for all n, $\mu_s \varphi_n^s \leq \mu_s^* \varphi_n^s$.

For a fixed s, let W_s and W_s^* represent the restrictions of V_s and V_s^* to $\mathcal{K}(\mathcal{P}(M \times \{h_\circ, h^\circ\}))$. Thus, for any $x \in \mathcal{K}(\mathcal{P}(M \times H))$, there exists $x' \in \mathcal{K}(\mathcal{P}(M \times \{h_\circ, h^\circ\}))$ such that $V_s(x) = W_s(x')$. Moreover, two DPF representations differ if and only if they differ on the domain $\mathcal{K}(\mathcal{P}(M \times \{h_\circ, h^\circ\}))$. Abusing notation again, let W_s and W_s^* denote the respective extensions to all compact subsets of aff $\mathcal{P}(M \times \{h_\circ, h^\circ\})$.

Recall that by Lemma 35, $V_s^*(y) \geq V_s(y)$ for all $y \in \mathcal{K}(\mathcal{P}(M \times H))$ and $V_s(h^\circ) > 0 > V_s(h_\circ)$. Then, by linearity, we must have $W_s^*(x) \geq W_s(x)$ for all compact $x \subset \operatorname{aff} \mathcal{P}(M \times \{h_\circ, h^\circ\})$. But $W_s(x_n) = \int_{\mathcal{U}} \varphi_n^s(u)\,\mathrm{d}\mu_s(u)$ (which implies, in particular, that each φ_n^s is integrable, since x_n is compact and $\max\{|W_s(x_n)|, |W_s^*(x_n)|\} < \infty$), so that we have $\mu_s\varphi_n^s \leq \mu_s^*\varphi_n^s$ as required.

"If": Consider the operators $\Phi, \Phi^* : C(H \times S) \to C(H \times S)$ defined as

$$\Phi W(f, s) := \sum_{s'} \Pi(s, s') \int_{\mathcal{U}} \max_{p \in x}\big[u(p_m) + \delta W(p_h)\big]\,\mathrm{d}\mu_s(u),$$

$$\Phi^* W(f, s) := \sum_{s'} \Pi(s, s') \int_{\mathcal{U}} \max_{p \in x}\big[u(p_m) + \delta W(p_h)\big]\,\mathrm{d}\mu^*(u).$$

As observed in the proof of Proposition 2, both Φ and Φ^* are monotone and also satisfy discounting, and are, therefore, contractions. For each $x \in \mathcal{K}(\mathcal{P}(M \times H))$, $\max_{p \in x}[u(p_m) + \delta W(p_h, s)]$ is a continuous and convex function of u. Therefore, by Theorem 5, for any $W \in C(H \times S)$, $\Phi W \leq \Phi^* W$. Let Φ^n and Φ^{*n} denote the nth iterates of Φ and Φ^*, respectively. We claim that $\Phi^n W \leq \Phi^{*n} W$ for all $W \in C(H \times S)$ and for all $n \geq 1$. We have already established this for $n = 1$. Suppose this is true for $n - 1$, that is, $\Phi^{n-1} W \leq \Phi^{*(n-1)} W$. Then $\Phi(\Phi^{n-1} W) \leq \Phi^*(\Phi^{n-1} W) \leq \Phi^*(\Phi^{*(n-1)} W)$, that is, $\Phi^n W \leq \Phi^{*n} W$, as claimed. Finally, let V and V^* be the unique fixed points of Φ and Φ^*, respectively, so that $V \leq V^*$, which completes the proof. *Q.E.D.*

PROOF OF PROPOSITION 9: Let U_s and $U_{s'}$, respectively, represent $\succsim_s$ and $\succsim_{s'}$ over $\mathcal{K}$, where, as before, $U_s(x) = \int_{\mathcal{U}} \max_{p \in x}[u(p_m) + \delta V(p_h, s)]\,\mathrm{d}\mu_s(u)$ and similarly for s'. Following the arguments in Lemma 35, we see that when restricted to consumption streams, $U_s|_L$ is a positive affine transformation of $U_{s'}|_L$. It is easy to see that the constant term must be zero, so let us suppose $U_s|_L = \lambda U_{s'}|_L$ for some $\lambda > 0$. Then it must be that μ_s and $\mu_{s'}$ differ by a scaling of λ so that $\mu_s u = \lambda \mu_{s'} u$. We also have $V_L(\cdot, s) = \lambda V_L(\cdot, s')$.

Now consider the S-IHCS $\ell_t^{\bar{m}}$ that delivers in each period, and in every state, p_m^*, the uniform lottery over M, except at time $t + 1$, where it delivers the prize $\bar{m} \in M$. Then $U_s(\ell_t^{\bar{m}})/U_{s'}(\ell_t^{\bar{m}}) = \lambda > 0$. Define $\eta_s := (0, \dots, 1, \dots, 0)$, where the 1 is the sth entry. The probability distribution over states S at time $t + 1$, conditional on being in state s at date 1 is $\eta_s \Pi^t$. Therefore, $U_s(\ell_t^{\bar{m}}) = \delta^t \sum_{\tilde{s}} \eta_s \Pi^t(\tilde{s}) \mu_{\tilde{s}} u_{\bar{m}}$, which implies that $\lambda = U_s(\ell_t^{\bar{m}})/U_{s'}(\ell_t^{\bar{m}}) = \frac{[\sum_{\tilde{s}} \eta_s \Pi^t(\tilde{s}) \mu_{\tilde{s}} u_{\bar{m}}]}{[\sum_{\tilde{s}} \eta_{s'} \Pi^t(\tilde{s}) \mu_{\tilde{s}} u_{\bar{m}}]}$ (or at least, such an $\bar{m} \in M$ can be chosen because $\mu_s u \neq \mathbf{0}$ for some $s \in S$). But Π is fully connected and has a unique stationary distribution π, which implies $\lim_{t \to \infty} \eta_s \Pi^t = \pi = \lim_{t \to \infty} \eta_{s'} \Pi^t$, which means we must have $\lambda = 1$, that is, $U_s|_L = U_{s'}|_L$, and, in particular, $\mu_s u = \mu_{s'} u$.

Consider now, the state $\tilde{s} \in S$ and let $\ell \in L$ be such that $V(\ell, \tilde{s}) \neq 0$. Such an ℓ exists because $V(\cdot, \tilde{s})$ is nontrivial over consumption streams. Let ℓ^* be the S-IHCS that gives the uniform lottery p_m^* in each state and in every period.

Now consider the consumption stream $\ell^{\dagger}$ that gives the consumption stream (p_m^*, ℓ) in state $\tilde{s}$ and the S-IHCS ℓ^* in every other state. Then $V(\ell^{\dagger}, s) = \Pi(s, \tilde{s})V(\ell, \tilde{s}) = \Pi(s', \tilde{s})V(\ell, \tilde{s}) = V(\ell^{\dagger}, s')$, which implies $\Pi(s, \tilde{s}) = \Pi(s', \tilde{s})$ because $V(\ell, \tilde{s}) \neq 0$.

Adapting the arguments from the proof of Theorem 2, we can now show the equivalence claimed, which proves the proposition. *Q.E.D.*

REFERENCES

AHN, D., AND T. SARVER (2013): "Preference for Flexibility and Random Choice," *Econometrica*, 81 (1), 341–361. [659]

ALIPRANTIS, C. D., AND K. C. BORDER (1999): *Infinite Dimensional Analysis: A Hitchhiker's Guide* (Second Ed.). New York: Springer. [666,678-680,686]

ANSCOMBE, F. J., AND R. J. AUMANN (1963): "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34, 199–205. [658,670]

BEKAERT, G., E. ENGSTROM, AND S. GRENADIER (2010): "Stock and Bond Returns With Moody Investors," *Journal of Empirical Finance*, 17 (5), 867–894. [660]

BLACKWELL, D. (1953): "Equivalent Comparison of Experiments," *Annals of Mathematical Statistics*, 24, 265–272. [700]

BOLLERSLEV, T., M. GIBSON, AND H. ZHOU (2011): "Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion From Option-Implied and Realized Volatilities," *Journal of Econometrics*, 160, 235–245. [659]

CAMPBELL, J., AND J. COCHRANE (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *The Journal of Political Economy*, 107 (2), 205–251. [660]

COCHRANE, J. H. (2005): *Asset Pricing*. Princeton, NJ: Princeton University Press. [668]

DEKEL, E., B. L. LIPMAN, AND A. RUSTICHINI (2001): "Representing Preferences With a Unique Subjective State Space," *Econometrica*, 69 (4), 891–934. [655]

DEKEL, E., B. L. LIPMAN, A. RUSTICHINI, AND T. SARVER (2007): "Representing Preferences With a Unique Subjective State Space: A Corrigendum," *Econometrica*, 75 (2), 591–600. [655]

DILLENBERGER, D., J. S. LLERAS, P. SADOWSKI, AND N. TAKEOKA (2013): "A Theory of Subjective Learning," Working Paper. [659]

FISHBURN, P. (1970): *Utility Theory for Decision Making*. New York: Wiley. [680]

GOLDMAN, S. M. (1974): "Flexibility and the Demand for Money," *Journal of Economic Theory*, 9, 203–222. [659]

GUL, F., AND W. PESENDORFER (2004): "Self-Control and the Theory of Consumption," *Econometrica*, 72 (1), 119–158. [656]

——— (2006): "Random Expected Utility," *Econometrica*, 74 (1), 121–146. [659]

HENDEL, I., AND A. NEVO (2006): "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, 74 (6), 1637–1673. [659,670]

HIGASHI, Y., K. HYOGO, AND N. TAKEOKA (2009): "Subjective Random Discounting and Intertemporal Choice," *Journal of Economic Theory*, 144 (3), 1015–1053. [658,664,668,675]

JONES, R. A., AND J. M. OSTROY (1984): "Flexibility and Uncertainty," *Review of Economic Studies*, 51 (1), 13–32. [659]

KOOPMANS, T. C. (1964): "On the Flexibility of Future Preferences," in *Human Judgments and Optimality*, ed. by M. W. Shelly and G. L. Bryan. New York: Wiley. [659,661]

KREPS, D. M. (1979): "A Representation Theorem for 'Preference for Flexibility'," *Econometrica*, 47 (3), 565–577. [655,656,658,663]

——— (1988): *Notes on the Theory of Choice*. Boulder, CO: Westview Press. [680]

KREPS, D. M., AND E. L. PORTEUS (1978): "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46 (1), 185–200. [661,664]

KRISHNA, R. V., AND P. SADOWSKI (2012): "Preference for Flexibility and the Pricing of Assets," Technical Report, Duke University. [660,676]

——— (2013): "Unobservable Persistent Taste Shocks," Technical Report, Duke University. [667,677]

LAX, P. D. (2007): *Linear Algebra and Its Applications*. Hoboken, NJ: Wiley. [695]

MAS-COLELL, A. (1978): "On Revealed Preference Analysis," *Review of Economic Studies*, 45 (1), 121–131. [673]

RUSTICHINI, A. (2002): "Preference for Flexibility in Infinite Horizon Problems," *Economic Theory*, 20, 677–702. [658]

SADOWSKI, P. (2013): "Contingent Preference for Flexibility: Eliciting Beliefs From Behavior," *Theoretical Economics*, 8, 503–534. [657,659]

SCHENONE, P. (2010): "Identifying Subjective Beliefs in Subjective State Space Models," Technical Report, Northwestern University. [659,664,683]

SMITH, D., AND R. WHITELAW (2009): "Time-Varying Risk Aversion and the Risk-Return Relation," Technical Report, New York University. [659]

TAKEOKA, N. (2006): "Representation With a Unique Subjective Decision Tree," Technical Report, Yokohama National University. [658]

TORGERSEN, E. (1991): *Comparison of Statistical Experiments*. Cambridge: Cambridge University Press. [700]

Dept. of Economics, Duke University, Durham, NC 27708, U.S.A.; vijay.krishna@duke.edu

and

Dept. of Economics, Duke University, Durham, NC 27708, U.S.A.; p.sadowski@duke.edu.

Manuscript received June, 2011; final revision received November, 2013.