REPEATED CHOICE: A THEORY OF STOCHASTIC INTERTEMPORAL PREFERENCES

JAY LU AND KOTA SAITO

Abstract. We provide a repeated-choice foundation for stochastic choice. We characterize when an agent's observed stochastic choice can be represented as a limit frequency of optimal choices over time. In our model, the agent repeatedly chooses today's consumption and tomorrow's continuation menu, aware that future preferences will evolve according to a subjective ergodic *utility process*. Using our model, we demonstrate how not taking into account the intertemporal structure of the problem may lead an analyst to biased estimates of risk preferences. A numerical example illustrates that these biases may be significant. Estimation of preferences can be performed by the analyst without explicitly modeling continuation problems (i.e. stochastic choice is *independent of continuation menus*) if and only if the utility process takes on the *standard* separable and linear form. Applications include dynamic discrete choice models when agents have non-trivial intertemporal preferences, such as Epstein-Zin preferences.

Keywords: Stochastic choice, intertemporal choice, dynamic choice, discrete choice, Epstein-Zin preferences, menu choice

Lu is affiliated with the Department of Economics, University of California, Los Angeles. Saito is affiliated with the Division of the Humanities and Social Sciences, California Institute of Technology. Lu's email address is jay@econ.ucla.edu; Saito's email address is saito@caltech.edu. We want to thank Andy Atkeson, Simon Board, Larry Epstein, Mira Frick, Faruk Gul, Ryota Iijima, Asen Kochov, Bart Lipman, Rosa Matzkin, Tomasz Strzalecki, John Rust, Yi Xing and seminar audiences at the SAET Conference at Academia Sinica, CIREQ Montreal Micro Theory Conference, Georgetown, Maryland, University of Montreal, Michigan, Rochester, University of Tokyo, LSE, Yale, LA Theory Conference, Caltech Junior Theory Workshop, RUD Conference at PSE, Zurich, Queen Mary, Royal Holloway, UCL, Pittsburg-Carnegie Mellon, Penn State, BU, UPenn and Ohio State for their helpful comments. We especially thank Kim Border for several discussions which helped us with the axiomatization. We also thank Jeffrey Zeidel for excellent research assistance for the numerical example. Financial support from the NSF under awards SES-1558757 (Saito), SES-1919263 (Saito) and SES-1919275 (Lu) are gratefully acknowledged.

1. Introduction

Modeling choice behavior as stochastic is common across many economic applications. In many of these applications, stochasticity is interpreted as a result of unobserved heterogeneity in a population of agents (henceforth, the "population interpretation"). On the other hand, the psychological origins of stochastic choice point to a single-agent interpretation.[1] There, stochasticity is interpreted as a result of a single agent making choices from the same decision problem repeatedly (henceforth, the "individual interpretation"). The literature on stochastic choice, however, has mostly taken such choice frequencies as given without considering when such a repeated-choice interpretation is possible and the underlying dynamic process generating stochastic choice.

In this paper, we provide the first theoretical foundation for this repeated-choice interpretation. Given an agent's stochastic choice, we obtain necessary and sufficient conditions under which the agent's observed stochastic choice can be represented as a limit frequency of his optimal repeated choices over time; in the representation, the agent repeatedly chooses today's consumption and tomorrow's continuation menu, aware that future preferences will evolve according to a *subjective utility process*.

Applying our model, we show that whenever the agent has non-trivial intertemporal preferences (e.g. Epstein-Zin preferences), his stochastic choice would be highly sensitive to the dynamic structure of future continuation menus. Slight changes to the structure (e.g. changing the frequency of how often choices are repeated) can affect stochastic choice in systematic ways. While we mostly focus on the single-agent interpretation of stochastic choice, our results also apply to the population interpretation and can be used to understand the systematic ways in which general intertemporal preferences affect the estimation of any dynamic discrete choice model.

We now provide an overview of our main results. First, we describe the formal setup. Based on the works of Kreps and Porteus (1978), Epstein and Zin (1989), and Gul and Pesendorfer (2004), we develop an infinite-horizon framework to study the agent's problem. Every period, the agent faces a *menu* (i.e., a choice set) which consists of risky prospects over consumption today and a continuation menu tomorrow. We focus on menus such that regardless of what the agent chooses or which outcome is realized, he will always face the same menu again after some finite time. Call such menus *repeated*. The structure of a repeated menu guarantees that the agent will choose from the same menu infinitely many

[1] Early work on models of stochastic choice include Thurstone (1927), Luce (1959), Block and Marschak (1960), and Falmagne (1978). The adoption of these models in economics to study unobserved heterogeneity naturally led to the population interpretation. For an overview of this history, see McFadden (2001).

times, generating a time series of choices. As a result, the agent's *stochastic choice* can be interpreted as the long-run frequency of choices from the repeated menu.[2]

Based on the setup, we introduce a new tractable model of stochastic choice. The agent's utility at time period t depends on some state variable s_t that evolves according to an ergodic Markov process. The Markov process is fixed and known to the agent but unknown to the analyst, which makes the agent's choice stochastic from the perspective of the analyst. For example, the state could be the agent's mood on a particular day, which affects how risk-averse and how impatient he is that day. Given the realization of state s_t at time t, the agent's utility of a pair (c, z) of today's consumption c and tomorrow's continuation menu z is recursively given by

$$u_t(c, z) = \phi_{s_t}\left(c, \mathbb{E}_{s_t}\left[\max_{p \in z} u_{t+1}(p)\right]\right). \tag{1}$$

There are two parts to this utility. First, the stochastic aggregator ϕ_{s_t} specifies the agent's intertemporal attitudes toward current consumption and future continuation value. Second, continuation values are evaluated by taking expectations with respect to the Markov process of the state. In other words, the agent is fully sophisticated; he knows the Markov process and takes expectations with the understanding that he will be choosing from the menu z tomorrow. The utility function (1) can be seen as a stochastic version of the model from Kreps and Porteus (1978) where continuation values are evaluated according to the linear representation of Dekel et al. (2001).

The *utility process* u_t defined in (1) is ergodic and describes the agent's stochastic intertemporal preferences at every time period t. In our representation theorem, for any menu z that repeats every t periods, the probability $\rho_z(p)$ that an option p is chosen from the menu repeated z is given by

$$\rho_z(p) = \lim_{n \to \infty} \frac{1}{n} \sum_{i=0}^{n} 1\left\{u_{it+1}(p) \geq u_{it+1}(q) \text{ for all } q \in z\right\}, \tag{2}$$

where $1\{\cdot\}$ is the indicator function and $u_t(p) = \int u_t(c, z)\, dp$ with $u_t(c, z)$ described as in (1). In this case, we say ρ is *ergodic*. Here, the probability that p is chosen from a set z is exactly the long-run frequency of the event that p is the best element in z according to the utility process. This is exactly the individual interpretation of stochastic choice models. We thus provide a theoretical foundation for this interpretation.

[2] We focus on repeated menus for simplicity and because of the fact that they are sufficient for identifying and characterizing our model, although we can extend our domain to include richer menus. For instance, we can consider menus that repeated with some probability. This extension is straightforward as long as the repetition probability does not depend on the agent's choice (otherwise, selection issues may complicate the identification exercise). See the discussion at the end of Section 2.1

The representation has strong uniqueness properties. Despite the generality of the model and the fact that our domain is restricted to only repeated menus, we show that the analyst can fully identify the agent's utility process from stochastic choice over repeated binary menus.

To study how intertemporal preferences affect estimation and inference in stochastic choice models, we present two applications. In the first application, we consider an analyst interested in eliciting an agent's risk aversion when the agent's stochastic aggregator ϕ takes on the well-known formula provided by Epstein and Zin (1989) (we call this *stochastic Epstein-Zin*). Understanding that the agent's preferences may be stochastic, the analyst asks the agent to repeatedly choose between a safe option (e.g., \$3 for sure) and a risky option (e.g., \$10 or \$0 with equal probability) repeatedly. We show that under stochastic Epstein-Zin, if his risk aversion is higher than his desire for consumption smoothing, then the probability of choosing the risky option increases when repetition becomes more frequent.[3] This is a novel behavioral phenomenon absent in stochastic choice models that do not explicitly address repetition.

We then show that if the analyst misspecified the model and ignored repetition, then she will *underestimate* the agent's atemporal risk aversion. To illustrate the magnitude of the bias, we consider a numerical example and find that even if the agent chooses the safe option in one time static choice, he will choose the risky option more than 80% of the time in a repeated choice setup. Thus, an analyst who ignores the intertemporal structure would incorrectly infer that the agent is mostly risk-loving. All this demonstrates the importance of modeling repetition when analyzing stochastic choice data.

In the second application, we consider a simple two-period example of a dynamic discrete choice model. Based on the same insight as in the first application, we illustrate the inherent inference issues that can arise if intertemporal preferences are not taken into account in applications of dynamic discrete choice estimation.

Both applications suggest that modeling repetition is crucial for inference when agents have more general intertemporal preferences. We also address the question of when an agent's preferences can be correctly inferred *without* modeling repetition explicitly. We define this formally using an axiom called *Independence of Continuation Menu (ICM)* and show that it is satisfied if and only if the utility process is *standard*, i.e., the stochastic aggregator takes the form of $\phi_s(c, v) = (1 - \beta_s) w_s(c) + \beta_s v$ where w_s is a random von Neumann–Morgenstern

[3] For such an agent, the risky option feels "safer" under repeated choice; intuitively, even if today's outcome is bad, repeating the choice means that there is always a chance that tomorrow's outcome will be good. As a result, the risky option becomes more attractive as repetition becomes more frequent.

(vNM) utility function and β_s is a random discount factor. In the case of stochastic Epstein-Zin preferences, *Indifference to Timing of Resolution of Uncertainty (IRU)* would ensure that the utility process is standard. In general, however, this is not true; we show that the gap between IRU and ICM is exactly a repeated version of the classic independence axiom, which we call *Repeated Independence (RI)*. We thus demonstrate the following three-way equivalence: ICM $\Leftrightarrow$ IRU + RI $\Leftrightarrow$ Standard Utility. The takeaway is that any generalization of standard utility will require the analyst to take into account repeated choice when conducting estimation or inference from stochastic choice. On the other hand, if agents are standard, then the analyst can conduct estimation ignoring continuation menus which would be useful in situations when continuation menus themselves may be unobservable.

Finally, we provide an axiomatic characterization of our model. While we focus only on the smaller domain of repeated menus, we show that the set of repeated menus is in fact dense in the set of all menus. In other words, for any generic menu z, we can construct a sequence of repeated menus that approximate z with arbitrary closeness. By considering a continuous extension, we can therefore focus on a stochastic choice function ρ over all finite (but not necessarily repeated) menus.

To axiomatize our representation, we construct a random expected utility model on an infinite-dimensional space where continuation menus are evaluated according to the representation in Dekel et al. (2001). This construction faces known technical challenges in the literature: extending Dekel et al. (2001) and Gul and Pesendorfer (2006) with countably-additive probability measures to an infinite-dimensional space has been considered to be difficult due to the lack of compactness in the infinite-dimensional setting.[4] We provide a unified methodology using the set of Lipschitz continuous utilities to address both challenges.

Our axioms combine the axioms of Gul and Pesendorfer (2006) with the linearity and continuity axioms of Dekel et al. (2001). We introduce three new axioms. The first two axioms (*Deterministic Stationarity* and *Average Stationarity*) are weaker analogs of the stationarity axiom of Koopmans (1960) for stochastic choice.[5] They allow us to construct a recursive and stationary Markov utility process. The last axiom (*D-continuity*) is a continuity condition stating that preference for flexibility is robust to small perturbations. It ensures ergodicity of the utility process. Finally, the representation is obtained by an application of the Birkhoff ergodic theorem. See the discussion after Theorem 4 for details.

[4] See Krishna and Sadowski (2014) for an outline of the technical issues of extending DLR. See Ma (2018) and Frick et al. (2018) for recent extensions of GP.

[5] A similar axiom appears in Lu and Saito (2018).

The rest of the paper is organized as follows. Section 2 introduces our repeated menus setup and our model with ergodic utilities. Section 3 presents the three applications of estimation under stochastic Epstein-Zin and dynamic discrete choice. In Section 4 we introduce ICM and its relationship with intertemporal preferences. Finally, Section 5 contains the axiomatic characterization. All omitted proofs are contained in the appendices.

1.1. **Related Literature.** Our paper is mainly related to four strands of literature in the following areas: (i) random expected utility, (ii) menu preferences, (iii) intertemporal choice, and (iv) dynamic discrete choice. The first strand of literature is on stochastic choice models of random expected utility. Gul and Pesendorfer (2006), Ahn and Sarver (2013), Lu (2016), and Lu and Saito (2018) study static models of stochastic choice, while Fudenberg and Strzalecki (2015) and Frick et al. (2018) study dynamic random choice.[6] Our paper is most closely related to the latter. The main differences are in motivation and the mathematical modeling. Given their motivation to study history dependency, Frick et al. (2018) study stochastic choice conditional on past menus, past choices, and consumption realizations, while our stochastic choice function is not conditional on these. Although they can interpret stochastic choice in their model as the result of a single agent, in contrast to our paper, they mainly focus on the population interpretation as it facilitates the interpretation of their primitive.[7] They consider any menus in a finite-horizon setup, while we consider repeated menus in an infinite-horizon setup.[8]

The second relevant strand consists of the modern literature on menu preferences, which began with Dekel et al. (2001) and Gul and Pesendorfer (2001). The former was extended to an objective state space by Dillenberger et al. (2014). Gul and Pesendorfer (2004) extends menu preferences to a dynamic setting by proposing an infinite-horizon consumption setup, which we have adopted in our paper. Other papers that make use of this framework include Higashi et al. (2009) and Krishna and Sadowski (2016). The first considers a random discounting model in which the agent anticipates the stochasticity of his future discount factor. The second extends the additive linear representation into an infinite-dimensional

[6] A more recent paper is Duraj (2018), which extends Frick et al. (2018) to a setting with an objective state space. Ke (2018) also studies expected utility in a Luce model.

[7] Given our motivation to provide a theoretical repeated-choice foundation for stochastic choice, we mainly focus on the individual interpretation although we can adopt the population interpretation (see Section 3.2).

[8] On the technical side, they also provide an extension of Gul and Pesendorfer (2006) to an infinite-dimensional setting. While they use the finiteness condition of Ahn and Sarver (2013) to extend the representation to a finitely additive measure, we use Lipschitz continuity to extend the representation to a countably additive one.

space. While their extension is finitely additive, our extension is countably additive while still preserving the uniqueness of the representation.[9]

Third, our paper is related to the classical literature on intertemporal choice. As mentioned, our model can be seen as a stochastic version of Kreps and Porteus (1978), including the popular special case of Epstein and Zin (1989) and Weil (1990). We also characterize a stochastic version of Uzawa-Epstein preferences, which was originally proposed by Uzawa (1968) and later axiomatized by Epstein (1983) in an extended setup with lotteries.[10] More recently, Bommier et al. (2017) also characterize standard utility via a monotonicity axiom.

Finally, our paper is related to the large literature on dynamic discrete choice. While the importance of considering more general intertemporal preferences (e.g., a preference for early resolution of uncertainty) is well-known, the literature has assumed standard additively separable preferences for the sake of tractability.[11] As far as we know, we are the first to analyze the effects of more general intertemporal preferences on inference under dynamic discrete choice. In addition, our ergodic representation (2) features in estimation methods of dynamic discrete choice models. Expanding on the work of Rust (1987), Hotz and Miller (1993) introduced an estimation methodology that is computationally less demanding. Their method of calculating conditional choice probabilities (CCP) from a sequences of choices uses a formula similar to our ergodic representation.[12]

2. A Model of Ergodic Utility

In this section, we first formally define repeated menus and then introduce our stochastic choice primitive. We then define a utility process and present our general model, an ergodic representation of stochastic choice. Finally, we discuss identification and uniqueness.

2.1. **Repeated Menus.** This section describes the basic setup of the model. Let time $T = \{1, 2, \dots\}$ be discrete and $M = [0, m]$ denote a closed interval representing consumption (e.g., money). The agent is faced with an infinite-horizon consumption problem (IHCP), that

[9] More recently, Krishna and Sadowski (2014) and Dillenberger et al. (2017) augment the dynamic setup with an informational structure. See Dillenberger et al. (2017) for a review of this literature.

[10] Recent papers that study the macroeconomic implications of stochastic intertemporal preferences include Alvarez and Atkeson (2017) and Barro et al. (2017).

[11] From Rust (1994), "expected-utility models imply that agents are indifferent about the timing of the resolution of uncertain events, whereas human decision-makers seem to have definite preferences over the time at which uncertainty is resolved. The justification for focusing on expected utility is that it remains the most tractable framework for modeling choice under uncertainty."

[12] On the other hand, a typical model in dynamic discrete choice assumes both observable states as well as unobservable states. While our model only includes unobservable states, it would be possible to extend our model to allow for observable states as well. This would allow for stochastic choices conditional on the observable state, which is exactly CCP.

is, a menu of choice options in which each option corresponds to a lottery over consumption today and a continuation menu tomorrow. We will refer to IHCPs simply as *menus* and denote them by $z \in Z$. From Gul and Pesendorfer (2004), we know that Z is homeomorphic to $\mathcal{K}(\Delta(M \times Z))$, where $\Delta(\cdot)$ denotes the set of probability measures and $\mathcal{K}(\cdot)$ denotes the set of nonempty compact subsets. Thus, we will associate Z with $\mathcal{K}(\Delta(M \times Z))$ without loss of generality.

Let $X = M \times Z$ denote the set of possible *outcomes*. For $x \in X$, we sometimes let $x \in \Delta X$ denote the degenerate lottery δ_x. For any $p \in \Delta X$, we let $p_M \in \Delta M$ and $p_Z \in \Delta Z$ denote its marginal distributions on M and Z respectively. We also sometimes use $p \in Z$ to denote the singleton menu $\{p\}$. Finally, let $ap + (1-a)q \in \Delta X$ denote the usual mixture between any two probability measures $p, q \in \Delta X$ and $a \in [0,1]$.

The main focus of our study will be on menus that repeat themselves after a fixed number of periods. The following example illustrates what we mean by such *repeated* menus.

Example 1 (Safe vs. Risky Option)**.** Consider an analyst interested in eliciting an agent's risk aversion which may be stochastic every period. Every day, the agent is offered a choice between a safe option b and a risky option r from the menu $z = \{b, r\}$. The safe option $b \in \Delta X$ yields \$3 for sure today and the menu $z \in Z$ again for sure tomorrow. The risky option $r \in \Delta X$ yields either \$10 or \$0 with equal probability today and the menu $z \in Z$ again for sure tomorrow. Note that the agent is sophisticated and understands that regardless of what he chooses today and which outcome is realized, he will always be faced with the menu z again for sure tomorrow. We will refer to this example in later sections.

Example 1 illustrates a menu that is repeated every period. More generally, we consider menus such that, regardless of what the agent chooses and which outcome is realized, he will always face the menu again for sure after a fixed number of time periods. Formally, for $z \in Z$, let $R_0(z) = \{z\}$ and for $t \in T$, define $R_t(z) := \mathcal{K}(\Delta(M \times R_{t-1}(z)))$. Thus, $R_t(z) \subset Z$ are the subset of menus that yield z for sure after t periods.

Definition. A menu z is *t-period* if $z \in R_t(z)$. The menu z is *repeated* if it is t-period for some $t > 0$.

The menu in Example 1 is 1-period since $z \in R_1(z)$. Let $Z^r \subset Z$ denote the set of repeated menus. In general, for a repeated menu, the agent will always face the *same* menu again after some fixed number of time periods. For example, if the menu is t-period, then the agent chooses from the menu at periods $1, 1+t, 1+2t$ and so forth. Since this is repeated ad infinitum, this can generate an infinite time series of choice data.

Repeated menus have several interesting properties. First, in a repeated menu, repetition is completely independent of the agent's choices. As a result, the analyst need not worry about selection biases interfering with the data collection process.

Second, even though repeated menus form a small subset of menus, they are rich enough to fully identify and characterize our model. In other words, the analyst can without loss only focus on repeated menus for identifying the parameters of our model (see Section 2.5). The reason for this is that repeated menus are dense in the set of all menus, i.e. they can be used to approximate any menu. We discuss this property further in Section 5.1.[13]

Finally, note that while we focus on repeated menus for simplicity and the fact that they are sufficient for identification and characterization, we can extend our setup to incorporate more common consumption-savings problems. For instance, instead of repeated menus that repeated with probability one, we can consider menus that repeated with some positive probability. For example, menu z could consist of two options, where each option yields z tomorrow with probability p and some other menu with probability $1-p$. Such extensions are straightforward as long as the repetition probability p does not depend on the agent's choice; otherwise, identification may be complicated by selection issues where a menu's occurrence depends on the agent's past choices. We can then approximate any consumption-saving problem by making the probability of repetition arbitrarily small.[14]

2.2. **Stochastic Choice.** In our model, the observable data, or primitive, is *stochastic choice*. Given repeated menus, we can interpret stochastic choice as the long-run frequency of the time series of choices. This interpretation of stochastic choice is standard in the literature, although it has not been modeled explicitly. For instance, in the random expected utility model of Gul and Pesendorfer (2006), stochastic choice can be interpreted as the long-run frequency of the time series choices from 1-period menus.[15] We now provide a formal definition of stochastic choice. Let $Z^f \subset Z$ denote the set of finite menus and let $Z^* = Z^r \cap Z^f$ denote the set of finite repeated menus.

Definition. A *stochastic choice* is a mapping $\rho : Z^* \to \Delta(\Delta X)$ such that for every $z \in Z^*$, ρ_z is a probability distribution on z.

[13] Moreover, it can be shown that there is always some minimal t^* for which z is t^*-period. Note that every t-period menu is also trivially kt-period for any positive integer k. In fact, t^* is the greatest common divisor of all possible periods of the menu; this is simply the first time z appears after the initial period.

[14] Note that if the menu never repeats, then choice is deterministic rather than stochastic choice as the analyst never observes the agent choosing again from the same menu.

[15] See Luce (1959) and Luce and Suppes (1965) for more detailed descriptions of the individual interpretation of stochastic choice.

Given a repeated menu $z \in Z^*$ and an option $p \in z$, the stochastic choice $\rho_z(p)$ designates the probability of choosing p from z. We deal with ties following Lu (2016) and Lu and Saito (2018) in allowing for some probabilities to be unspecified. This is analogous to how under standard deterministic choice, indifference characterizes exactly when the model is silent about which option the agent will choose. This approach allows the analyst to be agnostic about data that is orthogonal to the parameters of interest. For example, if two options are tied, then the stochastic choice is silent about the choice frequency for each option. Formally, we model this as non-measurability and let ρ denote the corresponding outer measure without loss of generality.[16] To simplify notation going forward, we sometimes let $\rho(z, y) = \rho_{z \cup y}(z)$ for $z, y \in Z^*$.[17] Thus, $\rho(p, q)$ denotes the frequency with which p is chosen over q.

We follow the literature on stochastic choice in assuming that the analyst only observes stochastic choice (i.e. the long-run choice frequency) and *not* the actual time path of choices. This is common in many empirical applications of stochastic choice, especially those that adopt the population interpretation. In dynamic discrete choice for instance, the analyst collects choices across both time and agents who are observationally identical under a standard i.i.d. assumption.[18] Since agents are i.i.d. across time, keeping track of the actual time series choice data is unnecessary so most models assume only stochastic choice is observable. For the individual interpretation, our paper is the first to connect stochastic choice with long-run choice frequencies; we represent stochastic choice as if it is generated from an infinite time series of optimal choices. Our focus on stochastic choice as a primitive is motivated by the existing literature and the fact that stochastic choice in our model is sufficient for identifying all the relevant parameters (see Theorem 1).[19]

2.3. **Utility Process.** In our model, the agent's utility at every period is stochastic and depends on the realization of state variable $s \in S$ that is unobserved by the analyst. We could interpret S as a set of subjective states that influence the agent's utility. For example, the state could be the agent's mood on a particular day, which affects how risk-averse or

[16] Let $\mathcal{F}$ be a σ-algebra on ΔX. Given any $z \in Z^*$, let ρ_z be a measure on the σ-algebra generated by $\mathcal{F} \cup \{z\}$. We can let ρ denote the outer measure with respect to this σ-algebra without loss of generality. See Lu (2016) for details.

[17] Note that if z contains no ties, then $\rho(z, y) = \sum_{p \in z} \rho_{z \cup y}(p)$ as all choice probabilities are specified. Otherwise, $\rho_{z \cup y}(z)$ denotes the outer measure.

[18] That is, the distribution of states is i.i.d. across both time and agents.

[19] Studying models that adopt time series choice data as a primitive would be interesting avenues for future research. In this case, the behavioral restrictions (on time series choice data) for representation would be more stringent as there will be different choice paths that generate the same long-run choice frequency. We thank Tomasz Strzalecki for discussions on this issue.

how patient he is on that day. We could also interpret the state as the realization of some private news arriving every period which affects the agent's utility that period.

The state evolves according to a Markov process $(s_t)_{t\in T}$ with transition probabilities $P : S \to \Delta S$ and a stationary distribution $\pi \in \Delta S$. The Markov process is fixed and known to the agent but unknown to the analyst. We assume the Markov process satisfies the continuity condition that $P_s \geq \delta\pi$ for some $\delta > 0$. This ensures that the Markov process has full support with respect to its stationary distribution and guarantees ergodicity.[20] Going forward, we let $[P]$ denote such a Markov process on the subjective state space.

We now describe the agent's utility. Let U denote the set of all utilities $u : X \to [0,1]$ normalized such that $u(\underline{x}) = 0$ and $u(\bar{x}) = 1$, where $\underline{x}$ and $\bar{x}$ correspond to consuming 0 and m forever, respectively.[21] For any measurable $u \in U$, we let $u(p) := \int_X u(x)dp$ denote the expected utility of $p \in \Delta X$.

Every period $t \in T$, a state $s_t \in S$ realizes and determines two things: (i) the agent's utility $u_{s_t} \in U$ at period t, and (ii) his expectation $\mathbb{E}_{s_t}$ about next period's state $s_{t+1} \in S$ according to the transition probability P_{s_t}. For example, the agent's mood determines his risk aversion and discount factor today and also informs his beliefs about his mood tomorrow. The agent is fully sophisticated and has correct beliefs; he anticipates what his mood will be tomorrow in order to determine his utility tomorrow as well as his beliefs about what his mood will be the day after, and so forth.

Following Kreps and Porteus (1978), we model utilities recursively as aggregator functions of current consumption and future continuation value. To accommodate changing utilities, we allow the aggregator function to be stochastic. A *stochastic aggregator* $\phi_s(c, v)$ specifies how the agent evaluates his current consumption c versus his future continuation value v given state $s \in S$. Formally, the stochastic aggregator $\phi_s : M \times [0,1] \to [0,1]$ is Lipschitz continuous (with some bound N) and strictly increasing in the second argument.[22] Since the agent anticipates that he may be choosing again next period, future continuation values are evaluated via the additive linear representation of Dekel et al. (2001). We now define a utility process as follows.

[20] The classic Doeblin's condition states that $P_s^n \geq \delta\lambda$ for some $n \geq 1$ and probability measure λ. Our condition obtains if we set $n = 1$ and $\lambda = \pi$.

[21] The range of utility need not be $[0,1]$ but any compact interval will suffice. Later, when we consider stochastic Epstein-Zin preferences, we use the range $[0,m]$ for convenience, where m is the largest monetary prize.

[22] Recall $[0,1]$ is the range of $u \in U$. If we change its range, then the domain of ϕ_s must be changed accordingly. For instance, when we consider stochastic Epstein-Zin preferences, ϕ_s is a function from $M \times M$ to M. All proofs and results go through.

Definition. A stochastic process $(u_t)_{t\in T}$ on U is a *utility process* if there exists a Markov process $[P]$ on S and a stochastic aggregator ϕ such that a.s.

$$u_t(c,z) = \phi_{s_t}\left(c, \mathbb{E}_{s_t}\left[\sup_{p\in z} u_{t+1}(p)\right]\right), \tag{3}$$

where the expectation $\mathbb{E}_{s_t}$ is taken with respect to P_{s_t}.

In this case, we say the utility process is *generated by* (P,ϕ). At a period $t\in T$, if $s_t = s$ for some $s\in S$, we sometimes write u_s or u_{s_t}, instead of u_t.

Every utility process is also an ergodic Markov process on the space of utilities. To see why it is a Markov process, note that if $u_s = u_{s'}$, then the agent's expectations $\mathbb{E}_s$ and $\mathbb{E}_{s'}$ are the same. Since the agent has correct beliefs, this means that the distribution of the next period's utility induced by P_s and $P_{s'}$ is also the same. Moreover, the following lemma shows that the utility process is ergodic as well. See Appendix A.1 for the proof.

Lemma 1. A utility process is an ergodic Markov process.

2.4. **Ergodic Representation of Stochastic Choice.** We are now ready to define the main model. We say the utility process is *regular* if $u_s(p) = u_s(q)$ with π-probability of either zero or one for all $p,q\in\Delta X$. In other words, ties either never occur or occur for sure.

Definition. ρ is *ergodic* if there exists a regular utility process generated by (P,ϕ) such that for any t-period $z\in Z^*$, a.s.

$$\rho_z(p) = \lim_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n} 1\left\{u_{it+1}(p)\geq u_{it+1}(q)\ \text{ for all } q\in z\right\},$$

If ρ is ergodic, then we say it is represented by (P,ϕ).

In our model, the stochastic choice of an option $p\in z$ corresponds exactly to the long-run frequency of optimally choosing p in an infinite sequence of choices by the agent. At every period, p is chosen only if it is ranked the highest in z according to realization of the utility process u. Recall that the utility process has a rich recursive intertemporal structure. Note that this is an as-if representation that corresponds exactly to the individual interpretation of stochastic choice in a repeated setup. Moreover, this features prominently in dynamic discrete choice estimation.[23] In Section 5, we provide an axiomatic characterization of the representation. For a simple illustration, consider a well-known special case of our model.

[23] For instance in Hotz and Miller (1993), similar formulas are used for the computation of conditional choice probabilities which are then used to estimate value functions for identifying parameters of interest. This methodological approach is now common in the literature.

Definition. A utility process is *standard* if there is a random vNM utility w_s and a random discount factor $\beta_s \in (0,1]$ such that a.s. $\phi_s(c,v) = (1-\beta_s) w_s(c) + \beta_s v$.

Standard utility is additively separable and linear in continuation values. It reduces to the random expected utility model of Gul and Pesendorfer (2006).

Example 2 (Random Expected Utility)**.** Let $[P]$ denote an i.i.d process and let the utility process to be standard. Thus, $u_t(c,z) = (1-\beta_s) w_{s_t}(c) + \beta_{s_t} \mathbb{E}\left[\sup_{p\in z} u_{t+1}(p)\right]$. Suppose ρ is represented by (P,ϕ). Consider a 1-period $z \in Z^*$. As a result, for any $p,q \in z$, we have $u_t(p) \geq u_t(q)$ if and only if $w_{s_t}(p_M) \geq w_{s_t}(q_M)$ from canceling out the continuation value of the menu z. From the ergodic representation, we thus have

$$\rho_z(p) = \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} 1\left\{w_{s_i}(p_M) \geq w_{s_i}(q_M)\, \forall q \in z\right\} = \pi\left\{s \in S : w_s(p_M) \geq w_s(q_M)\, \forall q \in z\right\},$$

where the second equality follows from the ergodic theorem. This corresponds exactly to the random expected utility model of Gul and Pesendorfer (2006).

Standard utility also corresponds to the classic Bellman equation in dynamic discrete choice. In that literature, an agent's utility satisfies $u_{s_t}(c,z) = (1-\beta)(w(c) + \varepsilon_{s_t}(c,z)) + \beta \mathbb{E}_{s_t}\left[\sup_{p\in z} u_{s_{t+1}}(p)\right]$, where the shocks ε are i.i.d. across both consumption c and continuation menus z. This coincides with standard utility when the continuation menu is the same across different choice options as in typical dynamic discrete choice problems (see Section 3.2).[24]

When the utility process is standard, the analyst does not need to model repetition explicitly. For instance, in Example 2, the agent's long-run choice frequency is independent of future continuation menus. However, this may no longer hold once we consider more general utilities processes, as we show formally in Section 4. For example, consider the widely-used intertemporal preferences of Epstein and Zin (1989) and Weil (1990). An interesting special case of our model which will feature prominently in our applications in Section 3 is stochastic Epstein-Zin.

[24] To see why, note that when the continuation menu is the same, we can suppress the dependency of the shocks on continuation menus. We can rewrite the shock $\varepsilon_s(c,z)$ simply as $\varepsilon_s(c)$, expressing the sum of current consumption utility w and the shock ε_s as a new current consumption utility $w_s(c) := w(c) + \varepsilon_s(c)$. As a result, the classic Bellman equation in dynamic discrete choice coincides with our standard utility model. One difference is that utilities in our model are bounded and Lipschitz continuous which would not be technically satisfied if shocks are extreme-value distributed. However, if we consider only a finite subset of choice options which is the case in most applications, then our conditions can be satisfied without loss of generality.

Definition. A utility process is *stochastic Epstein-Zin* if there are $RRA_s < 1$, $\psi_s < 1$, and $\beta_s \in (0,1)$ such that a.s. $\phi_s(c,v) = \left((1-\beta_s)\, c^{1-\psi_s} + \beta_s v^{\frac{1-\psi_s}{1-RRA_s}}\right)^{\frac{1-RRA_s}{1-\psi_s}}$.

In a stochastic Epstein-Zin utility process, each realized utility function is characterized by three stochastic parameters: (i) the relative risk aversion RRA,[25] (ii) the elasticity of intertemporal substitution $EIS = \psi^{-1}$, and (iii) the discount rate β. Since EIS captures how the agent is willing to shift consumption across periods in response to changes in interest rates, its reciprocal $\psi = EIS^{-1}$ can be interpreted as the agent's preference for consumption smoothing. Note that when $\psi = RRA$, this reduces to a special case of random expected utility with constant relative risk aversion (CRRA). If an ergodic ρ has a stochastic Epstein-Zin utility process, then we say ρ is *stochastic Epstein-Zin.*[26]

As in the deterministic case, there is a relationship between the Epstein-Zin parameters and preferences regarding the timing of the resolution of uncertainty (see Epstein et al. (2014) for a recent discussion). The following are extensions of the classic definitions of preference for early or late resolution of uncertainty in our repeated choice setup.

Definition. ρ satisfies *Preference for Early Resolution of Uncertainty (PEU)* if for all $\alpha \in [0,1]$, $\rho\left(\alpha\delta_{(c,z)} + (1-\alpha)\,\delta_{(c,y)}, \delta_{(c,\alpha z+(1-\alpha)y)}\right) = 1$; ρ satisfies *Preference for Late Resolution of Uncertainty (PLU)* if for all $\alpha \in [0,1]$, $\rho\left(\delta_{(c,\alpha z+(1-\alpha)y)}, \alpha\delta_{(c,z)} + (1-\alpha)\,\delta_{(c,y)}\right) = 1$.

It is well known that PEU corresponds to $\psi \leq RRA$ and PLU corresponds to $\psi \geq RRA$. As a result, indifference to the timing of the resolution of uncertainty (both PLU and PLR) corresponds to the case when $\psi = RRA$. This naturally extends to our setup as well (see Proposition 4 in Section 4).

2.5. **Identification and Uniqueness.** Given an ergodic representation, Theorem 1 below shows that the analyst can completely identify the agent's utility process from stochastic choice. In other words, the analyst does not require the full time series of choices for identification. Moreover, this can be done by focusing only on repeated binary menus.

Theorem 1. Let ρ and ρ' be represented by (P,ϕ) and (P',ϕ') respectively. Then $\rho(p,q) = \rho'(p,q)$ for all $p,q \in z \in Z^*$ if and only if (P,ϕ) and (P',ϕ') generate the same utility process.

[25] We follow Epstein and Zin (1989) in assuming $RRA < 1$ for simplicity. Technically when $RRA > 1$, the utility is decreasing in c but this can be ameliorated by dividing the utility by $1 - RRA$ as in typical applications of CRRA utility.

[26] A simple special case is when the subjective state space itself is $s = (RRA, EIS, \beta)$. Note that this is not without loss of generality, since utilities encode not only intertemporal preferences (in the form of the stochastic aggregator) but also the agent's expectations regarding tomorrow's state. The allowable subjective state space can thus be much richer than the three parameters (RRA, EIS, β).

See Appendix B for the proof. Given that stochastic choice data consist of only long-run frequencies, one may wonder how it would be possible to identify the agent's utility process completely beyond its stationary distribution. To see how this is possible, consider two different utility processes where one is i.i.d. and the other exhibits persistence but both have the same stationary distribution. Since the agent's utility also encodes information about his expectation regarding tomorrow's utilities, the analyst can distinguish between the two processes via the agent's attitudes toward continuation menus. Intuitively, in the i.i.d. case, tomorrow's utilities are more dispersed than in the persistent case so the agent would exhibit a greater preference for larger menus in the i.i.d. case than in the persistent case.[27]

3. Applications

We now demonstrate how intertemporal preferences affect estimation and inference in stochastic choice models. We present two applications. The first considers choice under risk with stochastic Epstein-Zin. The second considers a simple two-period dynamic discrete choice application.

3.1. **Choice Under Risk with Stochastic Epstein-Zin.** In this first application, we consider a setting where the agent chooses between risky and safe options. For instance, in Example 1, there is a safe option b that yields \$3 for sure and a risky option r that yields \$10 and \$0 with equal probability. In that example, repetition is modeled explicitly as occurring every period; in most models of stochastic choice (e.g., Gul and Pesendorfer (2006)) however, repetition is not modeled explicitly.

We first show how repetition affects the choice probability of the safe option for stochastic Epstein-Zin agents, introducing novel behavioral phenomena unique to stochastic choice. We then demonstrate a systematic way in which misspecifying intertemporal preferences bias risk aversion estimation. Finally, we provide a numerical example to demonstrate the magnitude of this bias.

3.1.1. *Delaying Repetition.* We first show how changing the frequency of repetition, e.g. delaying repetition, may significantly alter the agent's stochastic choice. This allows us to study novel behavioral phenomena resulting from the interaction between intertemporal preferences, risk and stochastic choice.

Given a 1-period menu z, consider delaying repetition so that the agent chooses every two periods instead of one. Call the new menu z^{+1}. In general, given a 1-period menu $z \in Z^*$ and

[27] Note that Theorem 1 does not mean that the Markov process on S can be identified uniquely; nonuniqueness can be easily obtained by relabeling or adding redundant states. Nevertheless, if we focus on a "minimal" state space such that no two states have the same utility, then unique identification holds.

$t \in T$, let z^{+t} denote the menu obtained by delaying repeated choice by t periods. Intuitively, for every $p \in z$, $p^{+t} \in z^{+t}$ can be expressed as follows: $p^{+t} = \left(p_M, \underbrace{0, \ldots, 0}_{t}; z^{+t}\right)$, where recall $p_M \in \Delta M$ is the marginal distribution of today's consumption. In other words, p^{+t} is the delayed lottery that yields the same marginal distribution of today's consumption as p but provides zero consumption for t-periods afterwards before repeating again. Note that z^{+t} is $t+1$-period. For an explicit example, consider Example 1. The delayed safe option b^{+1} will yield \$3 today, \$0 tomorrow, and z^{+1} the day after. The delayed risky option r^{+1} will yield \$10 and \$0 with equal probability today, \$0 tomorrow, and z^{+1} the day after. The delayed menu is $z^{+1} = \{b^{+1}, r^{+1}\}$.

The following result shows that when the agent's desire for consumption smoothing (i.e., ψ) is less (more) than risk aversion (i.e.,RRA), the probability that a safe option is chosen increases (resp., decreases) with delay.[28] These are novel behavioral phenomena unique to stochastic choice under risk.

Proposition 1. Suppose ρ is stochastic Epstein-Zin . Let z be a 1-period menu and $\delta_{(c,z)} \in z$ for some $c \in M$. Then, (i) $\psi_s \leq RRA_s$ a.s. implies $\rho_z\left(\delta_{(c,z)}\right) \leq \rho_{z^{+t}}\left((\delta_{(c,z)})^{+t}\right)$; moreover, $\rho_{z^{+t}}\left((\delta_{(c,z)})^{+t}\right)$ is increasing in t; (ii) $\psi_s \geq RRA_s$ a.s. implies $\rho_z\left(\delta_{(c,z)}\right) \geq \rho_{z^{+t}}\left((\delta_{(c,z)})^{+t}\right)$; moreover, $\rho_{z^{+t}}\left((\delta_{(c,z)})^{+t}\right)$ is decreasing in t.

Proof. Let $p = \delta_{(c,z)}$ where $z \in Z^*$ is 1-period and fix t_2 and t_1 such that $t_1 < t_2$. We will first show that if $\psi_s \leq RRA_s$ a.s., then $\rho_{z^{+t_1}}(p^{+t_1}) \leq \rho_{z^{+t_2}}(p^{+t_2})$. Consider the singleton continuation menu $r_i = (0, \ldots, 0; z^{+t_i})$ which gives 0 consumption for t_i periods. Let $v_i := \mathbb{E}_s\left[u_{\tilde{s}}(r_i)\right]$ denote its continuation value. Note that $v_1 \geq v_2$ as $t_1 < t_2$. Let $\sigma_s := \frac{1-\psi_s}{1-RRA_s}$ and since $\phi_s(c, v_2)^{\sigma_s} - \beta_s v_2^{\sigma_s} = (1-\beta_s) c^{1-\psi_s}$, we have $\phi_s(c, v_1) = \varphi_s(\phi_s(c, v_2))$, where $\varphi_s(x) := (x^{\sigma_s} + \beta_s(v_1^{\sigma_s} - v_2^{\sigma_s}))^{\sigma_s^{-1}}$. Note that $\varphi_s''(x) = (\sigma_2 - 1)(v_1^{\sigma_s} - v_2^{\sigma_s})\beta_s x^{\sigma_2 - 2}\varphi_s(x)^{1-2\sigma_s}$. If $RRA_s < 1$, then $\sigma_s \geq 1$ as $\psi_s \leq RRA_s$. On the other hand, if $RRA_s > 1$, then $\sigma_s < 0$ as $\psi_s < 1$. In either case, φ_s is convex so $\phi_s(\cdot, v_1)$ is more convex (i.e. more risk-averse) than $\phi_s(\cdot, v_2)$. For every $q \in z$, note that $q_M = q_M^{+t_1} = q_M^{+t_2}$ as the marginal distributions over today's consumption are the same. Thus, if $u_s(p^{+t_1}) = \phi_s(c, v_1) \geq \int_M \phi_s(c', v_1)\, dq_M = u_s(q^{+t_1})$, then $u_s(p^{+t_2}) = \phi_s(c, v_2) \geq \int_M \phi_s(c', v_2)\, dq_M = u_s(q^{+t_2})$, so (i) follows. The proof for (ii) where $\psi_s \geq RRA_s$ a.s. is analogous. □

To understand Proposition 1, recall Example 1 where the menu consists of a safe and and a risky option. Proposition 1 implies that whenever the agent's desire for consumption

[28] For convenience, we present Proposition 1 in its weak form but it also holds with strictness. That is, if $\psi_s < RRA_s$ holds with some probability, then $\rho_z\left(\delta_{(c,z)}\right) < \rho_{z^{+t}}\left((\delta_{(c,z)})^{+t}\right)$.

smoothing is smaller than his relative risk aversion (i.e., $\psi_s \leq RRA_s$), his probability of choosing the risky option increases when repetition becomes more frequent (i.e., the delay $+t$ becomes smaller). Intuitively, there are two forces at play. First, the risky option feels more attractive when repetition becomes more frequent because even if today's outcome is bad, there is always a chance that the agent will be lucky tomorrow and receive a good outcome. On the other hand, choosing the risky option will yield consumption that is more volatile over time. When preference for consumption smoothing is less than risk aversion, the first force dominates and the risky option is chosen more often when repetition is more frequent.

There are everyday examples where such behavior may be natural. For instance, a consumer may choose more "risky" brands if he knows he will visit the grocery store every day but stick to "safer" brands if he can visit the store only seldom.[29] A static model of stochastic choice that ignores repetition would fail to capture such behavioral phenomena. While we study delaying repetition as a simple example of changing continuation menus, our main point is that minor changes in continuation menus can profoundly alter stochastic choice in systematic ways.[30] We provide a more general analysis of such patterns in Section 4.

3.1.2. *Biased Estimation.* Suppose the analyst is interested in estimating the agent's risk aversion but misspecified the model and assumed standard instead of more general intertemporal preferences. This will naturally bias risk aversion estimates. We now show that there is a systematic way in which misspecification leads to biases.

Consider a stochastic Epstein-Zin agent as in Proposition 1 but $\psi_s = RRA_s$ a.s. In this case, the agent is standard with stochastic CRRA utility. As in Example 2, the probability that the agent will choose an option $p \in z$ is given by $\pi\{s \in S : w_s(p_M) \geq w_s(q_M) \,\forall q \in z\}$, where $w_s(p_M) = \int_M c^{1-RRA_s} dp_M$ is the expected CRRA utility of the marginal lottery p_M over today's consumption. Note that this coincides with the probability of choosing $p \in z$ in conventional stochastic choice models that are *static* and ignore the dynamic structure of repeated choice. As a result, using those models would lead to systematic biases in risk aversion estimates for general stochastic Epstein-Zin agents.[31]

[29] Another example is when people would be willing to bet on a repeated lottery but not on a one-time lottery as in the well-known Law of Large Numbers fallacy of Samuelson (1963).

[30] Proposition 1 assumes that when repetition is delayed, the agent receives zero (i.e., the lowest) consumption in the interim periods when there is no choice. Although this is a natural assumption, one may wonder what would happen if the agent receives the highest consumption in those periods. In this case, our results would naturally be flipped. That is, under stochastic Epstein-Zin, when the agent's risk aversion is higher (lower) than his desire for consumption smoothing, delaying repetition increases (decreases) his probability of choosing the safe option.

[31] As the case of Proposition 1, we present Proposition 2 in its weak form but it also holds with strictness.

Proposition 2. Suppose ρ is stochastic Epstein-Zin. Let z be a 1-period menu s.t. $\delta_{(c,z)} \in z$ for some $c \in M$. Then, (i) $\psi_s \leq RRA_s$ a.s. implies $\rho_z\left(\delta_{(c,z)}\right) \leq \pi\{s \in S : w_s(c) \geq w_s(q_M)$ for all $q \in z\}$; (ii) $\psi_s \geq RRA_s$ a.s. implies $\rho_z\left(\delta_{(c,z)}\right) \geq \pi\left\{s \in S : w_s(c) \geq w_s(q_M) \text{ for all } q \in z\right\}$.

Proof. Let $p \equiv \delta_{(c,z)}$ where $z \in Z^*$ is 1-period. We will first show (i) where $\psi_s \leq RRA_s$ a.s. By Proposition 1, we have $\rho_z(p) \leq \rho_{z^{+t}}(p^{+t})$ for any t. Consider the singleton continuation menu $r_t = (0, \ldots, 0; z^{+t})$, which gives 0 consumption for t periods. Let $v_t := \mathbb{E}_s[u_{\tilde{s}}(r_t)]$ denote its continuation value. Now, by making the delay t arbitrarily long, we can set the continuation value v_t arbitrarily small, i.e. $v_t \to 0$ as $t \to \infty$. Since $\phi_s(c, v) \to (1-\beta_s)^{\frac{1-RRA_s}{1-\psi_s}} w_s(c)$ as $v \to 0$, we have $\rho_z(p) \leq \lim_{t\to\infty} \rho_{z^{+t}}(p^{+t}) = \lim_{t\to\infty} \pi\{s \in S : \phi_s\left(c, v_t\right) \geq \int_M \phi_s\left(c', v_t\right) dq_M \forall q \in z\} = \pi\{s \in S : w_s(c) \geq w_s(q_M) \forall q \in z\}$, where the first equality follows from the ergodic theorem and the fact that $q_M = q_M^{+t}$ for all $q \in z$. Thus, (i) follows. The proof for (ii) where $\psi_s \geq RRA_s$ a.s. is analogous. □

Proposition 2 implies that an analyst who assumes the standard model would *underestimate* the agent's risk aversion if the agent is, in fact, stochastic Epstein-Zin with $\psi \leq RRA$. This is because the analyst would misinterpret the observed choice frequency $\rho_z(\delta_{(c,z)})$ as $\pi\left\{s \in S : w_s(c) \geq w_s(q_M) \text{ for all } q \in z\right\}$ and Proposition 2 shows the former is smaller than the latter when $\psi_s \leq RRA_s$ a.s. Thus, the analyst who observes a low probability of choosing the safe option would incorrectly infer that the agent's risk aversion is low. In reality, however, the agent is more willing to choose the safe option due to his intertemporal preferences.

That misspecification leads to biased estimates is unsurprising; the significance of Proposition 2 is that it demonstrates a systematic and tight pattern in the direction of bias. Under stochastic Epstein-Zin, if an agent's risk aversion is higher (resp., lower) than his desire for consumption smoothing, then ignoring the dynamic structure and intertemporal preferences would lead to underestimation (resp.,overestimation) of risk aversion. In the special case when $\psi = RRA$, there is no misspecification and Proposition 2 (i) and (ii) imply that there is no bias in estimating risk aversion.

Finally, while it may be possible to obtain a qualitative result similar to Proposition 2 while assuming deterministic Epstein-Zin preferences, the importance of our analysis is that it provides a framework to assess the *quantitative* magnitude of the bias in estimation. We show this explicitly in the numerical example below. To the best of our knowledge however, no such result–even a qualitative result–has appeared in the literature.[32]

[32] Epstein-Zin preferences have been widely used to resolve the equity premium puzzle in macro-finance. Those results, however, rely on equilibrium arguments that are intrinsically different from the analysis in Proposition 2.

3.1.3. *Numerical Result.* In actual empirical analysis, quantifying the magnitude of bias is important. This is useful especially for dynamic discrete-choice models in which agents receive preference shocks over time. We now provide an explicit numerical example to show that these biases can be highly significant and severe.

Consider Example 1 where every period, the safe option b yields \$3 for sure while the risky option r yields \$10 and \$0 with equal probability. Thus, $b = \delta_{(3,z)}$ and $r = \frac{1}{2}\delta_{(10,z)} + \frac{1}{2}\delta_{(0,z)}$ where $z = \{b, r\}$ is the 1-period menu. Let $r_M = \frac{1}{2}\delta_{10} + \frac{1}{2}\delta_0$ be the marginal distribution of consumption every period. The agent is stochastic Epstein-Zin and the state $s = (RRA, \psi)$ follows an i.i.d. process. We will assume the agent is so risk averse (i.e., $RRA \geq 0.5$ a.s) that, given a choice between a one-time payout of \$3 and a one-time 50-50 lottery of \$10 and \$0, he will choose the safe option.

Now we consider the agent's choice in the repeated menu z. From the Bellman equation, the continuation value of z is given by

$$v(z) = \int_S \max\left\{\phi_s(3, v(z)), \frac{1}{2}\phi_s(10, v(z)) + \frac{1}{2}\phi_s(0, v(z))\right\} ds. \tag{4}$$

Notice that since s is independently distributed across time, the continuation value $v(z)$ of the menu is the same across states. Thus the value of $v(z)$ can be obtained as a solution of equation (4).[33] The probability that the agent will choose the safe option b is then given by $\rho_z(b) = \pi\Big\{s \in S : \phi_s\big(3, v(z)\big) \geq \frac{1}{2}\phi_s\big(10, v(z)\big) + \frac{1}{2}\phi_s\big(0, v(z)\big)\Big\}$. The following result shows that this is significantly below unity.

Remark: Suppose that (i) ρ is stochastic Epstein-Zin with $\beta = 0.9$; (ii) ψ is uniform on $[0, 0.5]$; (iii) Independently of ψ, RRA is uniform on $[0.50, 0.97]$.[34] If $z = \{b, r\}$ is as in Example 1, then $\rho_z(b) \leq 20\% < 100\% = \pi\{s \in S : w_s(3) \geq w_s(r_M)\}$.

The magnitude of the bias is significant: the agent is risk-averse and will choose the safe option under static choice but will choose it less than 20% of the time under repeated choice due to intertemporal considerations. As a result, if the analyst either misspecifies the agent's preferences as standard or ignores the inherent dynamic structure of stochastic choice, then she will significantly underestimate risk aversion. Remark thus provides a numerical quantifier to statement (i) of Proposition 2.

[33] If s is not independently distributed across time, then the continuation value $v_s(z)$ of menu z should depend on s. We can calculate the value of $v_s(z)$ even for this case. In particular, when $\psi_s < RRA_s < 1$, it can be shown that $\phi_s(c, \cdot)$ is a contraction mapping which allows us to solve the Bellman equation for $v_s(z)$.

[34] We set the upper bound of RRA to 0.97 instead of 1 to avoid rounding errors due to computation (recall the fraction $1/(1 - RRA)$ appears in the utility function).

Our finding is robust across a wide range of parameter values. We varied the discount rate β from 0.70 to 0.99 and found that while the agent chooses the safe option under static choice, he will choose it under repeated choice at most 25% of the time.[35] We also varied the distribution of $s \equiv (RRA, \psi)$ by considering both binomial and beta distributions (assuming the independence between RRA and ψ). In all cases, the agent would choose the safe option under static choice but would choose it under repeated choice at most half the time.[36] Relaxing the independence of RRA and ψ also did not change our main findings.[37]

3.2. **Dynamic Discrete Choice.** In this second application, we apply our model to a simple two-period dynamic discrete choice example to illustrate the effects of intertemporal preferences on inference. The purpose of this application is to illustrate how our model can be readily applied to problems of discrete choice estimation that allow for more general temporal preferences.

Following most applications in dynamic discrete choice, we adopt the population interpretation of stochastic choice in this subsection only. In other words, we consider a population of observationally identical agents facing the same choice problem. This is possible in our model under two assumptions. First, even though choices are not technically repeated (we consider only two periods), we can model this as the limit of delaying repetition for an arbitrarily number of periods (see Section 3.1.1). Second, we assume the state follows an i.i.d. process where the distribution of each agent's state tomorrow is exactly equal to the population distribution π.[38] Under these assumptions, the long-run choice frequency that corresponds to stochastic choice also reflects the population choice. We can thus reinterpret stochastic choice in our ergodic model as a result of unobserved heterogeneity in a population of agents.[39]

The setup is as follows. There is a population of agents who decide whether to purchase phone insurance (e.g., AppleCare) at the beginning of years 1 and 2. We are interested in

[35] We tested β for $0.7, 0.8, 0.9$ and 0.99 and found $\rho_z(b)$ to be $24.8\%, 20.9\%, 17.8\%$ and 15.6% respectively. It is interesting that $\rho_z(b)$ is decreasing in the discount rate β. This is because as β increases, the continuation value also increases and $\phi_s(\cdot, v)$ becomes less concave. Note that the sensitivity of stochastic choice to the discount factor does not arise under standard intertemporal preferences.

[36] We tested beta distributions with parameter values $(2,2)$, $(0.5, 0.5)$, $(1,3)$ and $(3,1)$ and found $\rho_z(b)$ to be $8.9\%, 27.0\%, 0.5\%$ and 47.6%, respectively. For a binomial distributions with parameter 0.5, $\rho_z(b)$ is close to zero.

[37] We considered the case in which RRA and ψ are not independent; (RRA, ψ) is *jointly* uniformly distributed on the domain $0 \leq \psi < RRA \leq 0.97$. In this case, while the agent chooses the safe option 80% under static choice, he chooses it under repeated choice at most 32% of the time.

[38] We can relax this assumption as long as the stationary distribution of the (possibly non-i.i.d.) state process is the same as the population distribution.

[39] The latter assumption is a typical assumption when estimating conditional choice probabilities in the dynamic discrete choice literature. See Hotz and Miller (1993).

modeling their choice of insurance. Let c_s be the annual consumption value of the phone for an agent at state $s \in S$. We assume s is i.i.d. with stationary distribution π, which is also the population distribution of s. The price of insurance is a. In year $t \in \{1, 2\}$, there is p_t probability that the phone breaks down, in which case an agent's estimated cost for fixing a broken phone is θ_s. Both the consumer and the analyst know a, p_1, and p_2. Only the consumer knows the repair cost θ_s; the analyst would like to estimate the distribution of θ_s. For simplicity, we assume that $c_s \geq a$ and $c_s \geq \theta_s$ so all agents have positive final consumption. Note that in contrast to the application in the previous sections, utilities in this example appear stochastic to the analyst due to unobserved heterogeneity in the population (e.g., each agent's repair cost).

First, consider the case where all agents have risk-neutral standard preferences (i.e., stochastic Epstein-Zin from with $RAA_s = \psi_s = 0$). We study whether agents choose to buy insurance in year 1. Let β_s be the discount rate and v denote an agent's continuation value.[40] An agent will choose insurance if the following holds: $(1 - \beta_s)(c_s - a) + \beta_s v \geq p_1 ((1 - \beta_s)(c_s - \theta_s) + \beta_s v) + (1 - p_1)((1 - \beta_s)c_s + \beta_s v)$, or, equivalently, $\theta_s \geq a/p_1$. If we let b denote the "buy insurance" option, r denote the "not buy insurance" option and $z = \{b, r\}$ denote the menu, then the probability that insurance is purchased is given by $\rho^*_z(b) = \pi\{s \in S : \theta_s \geq a/p_1\}$. Naturally, lower values of θ_s correspond to fewer agents choosing insurance.

Next, we consider the case where all agents have more general intertemporal preferences. For instance, suppose the utility of an agent in state $s \in S$ is given by stochastic Epstein-Zin with risk neutrality (i.e., $RAA_s = 0$): $\phi_s(c, v) = \left((1 - \beta_s)c^{1-\psi_s} + \beta_s v^{1-\psi_s}\right)^{\frac{1}{1-\psi_s}}$, where ψ_s captures the agent's desire for consumption smoothing as in the previous subsection. Note when v is zero, this reduces to standard risk-neutral utility. The probability that insurance is chosen is given by $\rho_z(b) = \pi\{s \in S : \phi_s(c_s - a, v) \geq p_1 \phi_s(c_s - \theta_s, v) + (1 - p_1)\phi_s(c_s, v)\}$, where $v := \int_S \max\{\phi_s(b'), \phi_s(r')\} d\pi$ is the value of the continuation menu $z' = \{b', r'\}$, where b' and r' correspond to purchasing insurance or not respectively.

We now demonstrate how ignoring intertemporal preferences would lead to biased estimation of θ_s in this dynamic discrete choice problem. Suppose that agents' utilities are given by equation above, and, hence, the insurance adoption rate is given by $\rho_z(p)$. The analyst however misspecifies the model and assumes that utilities are standard. In this misspecified model, the insurance adoption rate is given by $\rho^*_z(p)$. The following proposition characterizes the comparison between $\rho^*_z(p)$ and $\rho_z(p)$ depending on the agents' intertemporal preferences.

[40] This is the same for all agents since the distribution of next period's state is π for everyone.

Proposition 3. (i) $\psi_s \leq 0$ a.s. implies $\rho_z(p) \leq \rho_z^*(p)$; (ii) $\psi_s \geq 0$ a.s. implies $\rho_z(p) \geq \rho_z^*(p)$.

Proof. Note that $\phi_s(\cdot, v)$ is convex if $\psi_s \leq 0$. Thus, $\phi_s(\cdot, v)$ is risk-loving so $\phi_s(c_s - a, v) \geq p_1 \phi_s(c_s - \theta_s, v) + (1 - p_1)\phi_s(c_s, v)$ implies $c_s - a \geq p_1(c_s - \theta_s) + (1 - p_1)c_s$. This means that $\rho_z(p) \leq \rho_z^*(p)$ as desired. The case for $\psi_s \geq 0$ is symmetric. □

Proposition 3 implies that if ψ_s is negative for almost all agents, then ignoring intertemporal preferences will result in *underestimation* of repair costs.[41] To see this, note that the analyst misinterprets the observed adoption rate $\rho_z(p)$ as $\rho_z^*(p)$ and will estimate θ_s based on the misspecified model. Proposition 3 shows that $\rho_z(p) \leq \rho_z^*(p)$ when ψ_s is negative a.s. This means that if the analyst observes a low adoption rate, she would incorrectly infer that repair costs are low.[42] In reality however, agents are more willing to decline insurance due to their intertemporal preferences. The implication for when ψ_s is positive for almost all agents is symmetric.[43]

This example illustrates how our model can be readily applied to problems of discrete choice estimation that allow for more general temporal preferences. Although we assumed risk neutrality for simplicity, this example can be easily generalized to accommodate non-trivial risk attitudes. Our example is straightforward but it serves to illustrate the inherent inference issues that can arise if intertemporal preferences are not taken into account in many applications of dynamic discrete choice estimation. While ignoring intertemporal preferences would obviously affect inference, our main point is understanding the systematic way in which intertemporal preferences affect estimation as outlined in Proposition 3.[44]

4. Independence of Continuation Menus

In Section 3, we demonstrated how the explicit modeling of repeated choice is paramount for an analyst interested in elicitation or inference when the agent has more general intertemporal preferences. In this section, in a more general setup, we formalize when repeated choice

[41] Since we are considering risk-neutral agents ($RRA_s = 0$), ψ_s is negative corresponds to preference for early resolution of uncertainty.

[42] Recall that a lower adoption rate corresponds to lower values of θ_s.

[43] The intuition for Proposition 3 parallels that of Proposition 2. Note that buying (not buying) insurance in Proposition 3 corresponds to choosing the safe option (resp., the risky option) in Proposition 2. This is because if agents purchase insurance, their payoffs are constant. Therefore, under the assumption of risk neutrality (i.e., $RAA = 0$), statements (i) and (ii) in Proposition 3 correspond respectively to statements (i) and (ii) in Proposition 2.

[44] It is easy to incorporate additive shocks widely used in the dynamic discrete choice literature. With additive errors ε_s, the utility if the phone breaks is given by $u_s(c_s - \theta_s, z) = \phi_s(c_s - \theta_s + \varepsilon_s(c_s - \theta_s), v_s(z))$, where θ_s is the repair cost. The same argument for Proposition 3 then applies in this setting. If ψ_s is negative for almost all agents, then ignoring intertemporal preferences will result in underestimation of repair costs. Vice-versa, if ψ_s is positive for almost all agents, then ignoring intertemporal preferences will result in overestimation of repair costs.

needs to be taken into account by the analyst versus when it is unnecessary to do so as in static random choice. In the case of the latter, we say the stochastic choice satisfies an axiom called *Independence of Continuation Menus.*

To illustrate, recall Example 1 where the menu consists of a risky option that yields \$10 and \$0 with equal probability and a safe option that yields \$3 for sure. Proposition 1 implies that the probability of choosing the risky option over the safe option depends on the timing of the next repetition; in other words, continuation menus matter unless the agent is indifferent to the timing of resolution of uncertainty. On the other hand, in Example 2 where we assume standard utility, the only thing that matters is the distribution of current consumption; in that case, choice is independent of continuation menus.

We now formalize these concepts. Fix a menu $z \in Z$. Now, for any $p \in z$, recall that $p_M \in \Delta M$ and $p_Z \in \Delta Z$ are the marginal distributions of today's consumption and tomorrow's continuation menu respectively. Let $z_M := \{p_M \in \Delta M : p \in z\}$ denote the menu of distributions of today's consumption. Note that if $p_Z = q_Z$ for all $p, q \in z$, then the distribution of tomorrow's continuation menu is independent of what the agent chooses. We call such a menu 1-*period invariant*.[45]

The following definition characterizes when choice is independent of tomorrow's continuation menus. Consider a menu z where $p_Z = r$ for all $p \in z$ so z is 1-period invariant. Now, construct another menu from z by switching the distribution of tomorrow's menus from r to r' but leaving the distribution of today's consumption the same. Call this new menu y. In other words, $z_M = y_M$ and $q_Z = r'$ for all $q \in y$. Note that both z and y are 1-period invariant. *1-Period Independence of Continuation Menus* states that choice probabilities in y and z are the same; in other words, switching the common distribution of tomorrow's menus does not alter stochastic choice.

Definition. ρ satisfies *1-Period Independence of Continuation Menus (1-ICM)* if for all 1-period invariant $z, y \in Z^*$, $p \in z$ and $q \in y$,

$$p_M = q_M \text{ and } z_M = y_M \implies \rho_z(p) = \rho_y(q).$$

Under 1-ICM, the agent evaluates today's consumption independent of tomorrow's continuation menus.[46] This follows from the fact when today's consumption is evaluated independent of tomorrow's continuation menus, the agent will naturally ignore correlations between today's consumption and tomorrow's menus.

[45] Note that every 1-period menu is 1-period invariant. The converse is not true.

[46] In fact, it implies the separability axiom of Frick et al. (2018) which is the stochastic analog of the standard separability axiom of Fishburn (1970).

1-ICM is applicable only to menus that are 1-period invariant. This is the case in Proposition 1 where z is 1 period and $y = z^{+t}$ for some t so $z_M = y_M$.[47] In general however, we may consider menus that are *not* 1-period invariant.[48] We now extend our notion of independence beyond the first period. For simplicity, we will focus on menus such that every continuation menu before time t is degenerate. We call such menus *t-simple*. For every option in a t-simple menu, we can consider its distributions over t-period consumptions and continuation menus. Formally, let $M_1 := M$, and recursively define $M_t := M \times \Delta M_{t-1}$. Let $p_{M_t} \in \Delta M_t$ denote the t-period distribution of consumption and let $z_{M_t} := \{p_{M_t} \in \Delta M_t : p \in z\}$ denote the menu of t-period consumption distributions. Also let $p_{Z_t} \in \Delta(\Delta(\cdots \Delta Z))$ denote the t-period distribution of continuation menus where the $\Delta(\cdot)$ operator is applied t times. Given a menu z, if $p_{Z_t} = q_{Z_t}$ for all $p, q \in z$, then the menu is *t-period invariant*.[49]

The next definition characterizes when choice is independent of all continuation menus. It extends 1-ICM from 1 period to t periods. Similar to the reasoning for 1-ICM, ICM implies that switching the common distribution of continuation menus does not alter stochastic choice.

Definition. ρ satisfies *t-Period Independence of Continuation Menus (t-ICM)* if for all t-period invariant $z, y \in Z^*$, $p \in z$ and $q \in y$,

$$p_{M_t} = q_{M_t} \text{ and } z_{M_t} = y_{M_t} \implies \rho_z(p) = \rho_y(q).$$

ρ satisfies *Independence of Continuation Menus (ICM)* if it satisfies t-ICM for all $t \in T$.

In the following, we characterize utility processes that satisfy ICM. To state the main result of this section, we introduce a definition: A utility process is *additively separable* if there is a random vNM utility w_s, a random function φ_s and a random discount factor $\beta_s \in (0,1]$ such that a.s. $\phi_s(c, v) = (1-\beta_s) w_s(c) + \beta_s \varphi_s(v)$. Note that an additively separable utility process is more general than the standard model: it is standard if and only if $\varphi_s(v) = v$ a.s.

Theorem 2. Suppose ρ is ergodic. Then, (i) ρ satisfies ICM if and only if its utility process is standard; (ii) ρ satisfies 1-ICM if and only if its utility process is additively separable; (iii) ρ satisfies ICM if and only if it satisfies 1-ICM and 2-ICM.

[47] In Proposition 1, note that for all $p \in z$, $p_Z = \delta_z$ while for all $q \in y$, $q_Z = \delta_{(0,\dots,y)}$ which corresponds to 0 consumption for t periods followed by y. Hence, both y and z are 1-period invariant. Proposition 1 implies that if $\psi_s = RRA_s$ a.s., then $\rho_z\left(\delta_{(c,z)}\right) = \rho_{z^t}\left(\delta^{+t}_{(c,z)}\right)$ which agrees exactly with 1-ICM.

[48] Suppose the analyst is interested in eliciting the agent's discount factor. In order to do this, she would need to offer repeated menus of *at least* 2 periods.

[49] As in 1-period menus, every simple t-period menu is t-period invariant but the converse is not true.

See Appendix C.1 for the proof. Theorem 2 (i) demonstrates the equivalence between ICM and standard utility.[50] While it may not be surprising that standard utility ensures ICM, it is notable that ICM implies standard utility. In other words, whenever the agent has non-standard intertemporal preferences (i.e., non-standard utility), there exists some repeated choice problem where continuation menus matter; ignoring repeated choice in such a problem would result in biased inference. Recall that standard utility was exactly the case assumed in Example 2 where our ergodic model reduces to the static model of random expected utility.

Theorem 2 (ii) and (iii) show that while the additive separability is sufficient to ensure 1-ICM, it is insufficient to ensure ICM. In other words, when an agent has an additively separable utility process, the analyst can ignore repetition for 1-period menus but not 2-period ones.

4.1. **Repeated Independence.** In this subsection, we relate ICM with other well-studied behavioral properties. We show that while indifference to the timing of resolution of uncertainty is not sufficient to characterize ICM, adding a *repeated version* of the independence axiom is enough. This serves two purposes. First, it allows us to clarify the relationship of ICM with other common behavior studied in the literature. Second, given the equivalence between the standard utility model and ICM (see Theorem 2), this provides an alternate characterization of the standard utility model commonly used in dynamic discrete choice.

Consider the stochastic Epstein-Zin preferences of Section 3.1 and note that if the agent satisfies *Indifference to Timing of Resolution of Uncertainty (IRU)* (i.e., both PEU and PLU), then the utility process is standard (i.e., $\psi_s = RAA_s$ a.s.). Given Theorem 2, this means that under stochastic Epstein-Zin preferences, IRU ensures that ICM is satisfied. For general utility processes however, IRU does not imply ICM. A utility process is *stochastic Uzawa-Epstein* if there are vNM utilities w_s and $\beta_s(\cdot) \in (0,1]$ such that a.s. $\phi_s(c,v) = (1-\beta_s(c))\, w_s(c) + \beta_s(c)\, v$.

Proposition 4. Let ρ be ergodic. (i) ρ satisfies PEU (PLU) if and only if $\phi_s(c,\cdot)$ is convex (resp., concave) a.s.; (ii) it satisfies IRU if and only if its utility process is stochastic Uzawa-Epstein.

Proof. Suppose ρ is ergodic and exhibits PEU. Then a.s. $\alpha\phi_s(c, v_s(z)) + (1-\alpha)\,\phi_s(c, v_s(y)) \geq \phi_s(c, \alpha v_s(z) + (1-\alpha)\, v_s(y))$. Since this is true for all z and y, the result follows. The

[50] It corresponds to an infinite-horizon Markovian version of the Bayesian Evolving Utility model of Frick et al. (2018).

case for PLU is symmetric. If $\phi(c,\cdot)$ is both concave and convex, then it is linear. Thus, $\phi_s(c,v) = (1-\beta_s(c))\, w_s(c) + \beta_s(c)\, v$ for $\beta_s(c) > 0$ for all $c \in M$. □

Proposition 4 is the stochastic analog of Theorem 1 of Epstein (1983). Since the discount factor β depends on consumption, this representation is more general than the standard model. In fact, since Uzawa-Epstein utilities are not additively separable, it follows from Theorem 2 (ii) that IRU will not even ensure 1-ICM.

Given that IRU does not ensure ICM but ICM implies IRU (since every standard utility satisfies IRU), it is natural to ask what additional property will bridge the gap between IRU and ICM. It turns out to be a repeated version of classic independence axiom. To illustrate, recall Example 1 where the 1-period menu z consists of a risky option that yields \$10 and \$0 with equal probability and a safe option that yields \$3 for sure. Suppose we wanted to test the independence axiom in this repeated setup by mixing both the risky and safe options with a third option r that yields \$4 for sure. Let y denote this new 50-50 mixture of z and r. Note that y is also a 1-period menu and consists of two options: one option that yields \$10 with probability 0.25, \$0 with probability 0.25, and \$4 with probability 0.50; the other option yields \$3 and \$4 with equal chance. Importantly, regardless of what happens, the agent will face y for sure next period so this mixture is repeated every period ad infinitum. We use the notation $y = 0.5z \otimes 0.5r$ to denote this 50-50 repeated mixture between z and r. This corresponds exactly to repeated testing of the classic independence axiom.

We now formalize this concept. First consider a 1-period menu $z \in Z^*$ in which every $p \in z$ can be expressed as $(p_M; z)$. Consider repeatedly mixing z with some $r \in \Delta M$. This yields the new 1-period menu, denoted by $\alpha z \otimes (1-\alpha)\, r \in Z^*$, such than any element of the 1-period menu is of the form $(\alpha p_M + (1-\alpha)\, r; \alpha z \otimes (1-\alpha)\, r)$. In other words, every option is mixed with r every period. We denote the element of $\alpha z \otimes (1-\alpha)\, r \in Z^*$ by $\alpha p \otimes (1-\alpha)\, r$. We can extend this to all t-period simple menus (see Appendix S.4) and define repeated independence as follows.

Definition. ρ satisfies *Repeated Independence (RI)* if for all t-simple $z \in Z^*$, $\alpha \in (0,1)$ and $r \in \Delta M$, $\rho_z(p) = \rho_{\alpha z \otimes (1-\alpha) r}(\alpha p \otimes (1-\alpha)\, r)$.

RI is exactly the classic independence axiom in our repeated choice setup. In fact, it corresponds to the linearity axiom in the static random expected utility model of Gul and Pesendorfer (2006). The main result of this subsection shows that IRU in addition to RI exactly characterizes ICM. Moreover, under IRU, RI is equivalent to 1-ICM.

Theorem 3. Suppose ρ is ergodic. Then the following statements are equivalent: (i) ρ satisfies ICM; (ii) ρ satisfies IRU and RI; (iii) ρ satisfies IRU and 1-ICM; (iv) the utility process is standard.

Proof. For the equivalence between (i) and (ii), see Appendix C.3. The equivalence between (i) and (iii) follows from Theorem 2, Proposition 4, and the fact that any additively separable Uzawa-Epstein utility must be standard. The equivalence between (i) and (iv) follows from Theorem 2. □

The main result of Theorem 3 is the equivalence between (i) and (ii): IRU together with RI exactly characterizes ICM. The equivalence between (ii) and (iv) also has important implications in the dynamic discrete choice literature given the widespread use of standard utility. As mentioned, Rust (1994) pointed out that preferences for early or late resolution of uncertainty have been ignored in that literature as standard utility implies IRU (see footnote 11). The equivalence between (ii) and (iv), however, shows that RI is yet another implication of standard utility.[51]

5. Characterization

This section provides an axiomatic characterization of our model. First, we show how repeated menus can be used to approximate any menu. This allows us to extend our primitive to the set of all (finite) menus.

5.1. **Extending Repeated Menus.** Given any menu $z \in Z$, consider replicating the menu z for the first t periods and ending with a menu $y \in Z$ for sure. We use the notation $r_{y,t}(z)$ to denote such a menu and construct it inductively as follows. First, for any $y \in Z$, let $r_{y,0}(z) = y$. Given $r_{y,t-1}$, for any $p \in \Delta X$, let $p_{y,t} \in \Delta X$ denote the lottery induced by $r_{y,t-1}$, that is, for all measurable $A \times B$, $p_{y,t}(A \times B) = p\left(A \times r_{y,t-1}^{-1}(B)\right)$. Finally, for any $z \in Z$, define $r_{y,t}(z) := \{p_{y,t} : p \in z\}$. In other words, $r_{y,t}(z) \in Z$ is the menu that follows z for the first t periods ending with y for sure. Lemma 10 in Appendix S.1 shows that this is well-defined. Given any menu $z \in Z$, we can now define what it means to construct a repeated menu that approximates z up to t periods. We let z^t denote this t-period repeated version of z.

[51] Theorem 3 also suggests that with stochastic choice, intertemporal preferences complicate tests of the classic independence axiom. Even though the agent may satisfy the static independence axiom for a single time period, he may violate this repeated version of the independence axiom (i.e., RI). In Appendix S.4, we study the relationship between more general intertemporal preferences and particular patterns of RI violations along with comparative statics. Moreover, as we will show in the next section, any ergodic ρ satisfies the independence axiom over menus (i.e., Linearity, Axiom 1.2). These facts show the importance of specifying the appropriate domain when we test the independence axiom with stochastic choice.

Definition. Given $z \in Z$, let z^t be t-period such that $z^t = r_{z^t,t}(z)$.

The following lemma shows this is well-defined. Moreover, given any menu $z \in Z$, we can use its t-period repeated version to approximate it as we increase the number of periods between each repetition. This, in turn, implies that we can use finite repeated menus to approximate any finite menu in Z^f. See Appendix S.1 for the proof.

Lemma 2. (i) For every $z \in Z$, z^t exists and $z^t \to z$ as $t \to \infty$; (ii) Z^* is dense in Z^f.

5.2. **Axiomatic Characterization.** The results in the previous section allow us to extend the observed stochastic choice on repeated finite menus to all finite menus as follows. Consider a random choice $\bar{\rho}$ on all finite menus Z^f such that $\bar{\rho}_z = \rho_z$ for every $z \in Z^*$. In other words, $\bar{\rho}$ agrees with ρ on all repeated menus Z^*. From Lemma 2 (ii), we know that Z^* is dense in Z^f. Thus, for any $z \in Z^f$, we can find $z^t \in Z^*$ such that $z^t \to z$. If $\bar{\rho}$ is continuous, then ignoring ties, $\bar{\rho}_z = \lim_t \rho_{z^t}$. Thus, we can think of $\bar{\rho}$ as the continuous extension of ρ from Z^* to Z^f. We model ties in the same way as ρ (see the discussion on ties in Section 2.1) and let $Z^\circ \subset Z^f$ denote the set of finite menus that contain no ties. To simplify notation going forward, we let ρ denote $\bar{\rho}$ without loss of generality.

We are now ready to state our axioms on stochastic choice. The first set of axioms consists of conditions on random expected utility. Note that mixtures here are taken ex-ante at time 0 and we let $\text{ext}(z)$ denote the extreme options of some menu $z \in Z^f$. Also recall that $\bar{x}$ and $\underline{x}$ are the consumption streams that yield the best outcome (i.e., m) and worst outcome (i.e., 0) respectively forever. We sometimes let x denote the singleton menu that yields consumption $x \in X$ forever.

Axiom 1.1 (Monotonicity)**.** If $z \subset y$, then $\rho_z(p) \geq \rho_y(p)$ for any $z, y \in Z^f$ and $p \in z$,

Axiom 1.2 (Linearity)**.** $\rho_z(p) = \rho_{\alpha z + (1-\alpha) q}(\alpha p + (1-\alpha) q)$ for any $p \in z \in Z^f$, $\alpha \in (0,1)$ and $q \in \Delta X$.

Axiom 1.3 (Extremeness)**.** $\rho_z(\text{ext}(z)) = 1$ for any $z \in Z^f$.

Axiom 1.4 (Continuity)**.** $\rho : Z^\circ \to \Delta(\Delta X)$ is continuous.

Axiom 1.5 (Best-Worst)**.** $\rho(\underline{x}, \bar{x}) = 0$ and $\rho(\bar{x}, x) = \rho(x, \underline{x}) = 1$ for all $x \in X$.

Axiom 1.6 (L-continuity)**.** There exists $N > 0$ such that for any $\alpha \in [0,1]$ and any $x, x' \in X$, $|x - x'| \leq \frac{\alpha}{N}$ implies $\rho\left(\alpha \delta_{\bar{x}} + (1-\alpha)\delta_x, \alpha \delta_{\underline{x}} + (1-\alpha)\delta_{x'}\right) = 1$.

Axioms 1.1-1.4 are direct from Gul and Pesendorfer (2006). Best-Worst (Axiom 1.5) ensures that $\bar{x}$ and $\underline{x}$ truly are the best and worst outcomes. Finally, L-continuity (Axiom

1.6) is the stochastic version of the Lipschitz continuity axiom from Dekel et al. (2007). It guarantees that utilities are sufficiently well-behaved in that they are Lipschitz continuous with respect to some common bound N. This is important for the representation and ensures that it is unique.[52] To understand L-continuity intuitively, note that when $N = \infty$, then $x_1 = x_2 = x$ for some x and the axiom reduces to $\rho\left(\alpha\delta_{\bar{x}} + (1-\alpha)\,\delta_x, \alpha\delta_{\underline{x}} + (1-\alpha)\,\delta_x\right) = 1$, which holds by Best-Worst and Linearity. L-continuity requires that this holds for large enough but finite N.[53] Taken together, Axiom 1 characterizes a random expected Lipschitz utility with best and worst outcomes.

Continuation Linearity (Axiom 2) below ensures that agent's preference toward continuation menus satisfy linearity with respect to ex-post mixing. First, we define component-wise ex-post mixing. For $\lambda \in [0,1]$, $c, c' \in M$ and $z, z' \in Z$, define ex-post mixing as $\lambda\delta_{(c,z)} \oplus (1-\lambda)\,\delta_{(c',z')} := \delta_{(\lambda c+(1-\lambda)c', \lambda z+(1-\lambda)z')}$. Here, the first mixture $\lambda c + (1-\lambda)\,c'$ corresponds to the standard mixing of monetary consumptions (i.e., real numbers) while the second mixture $\lambda z + (1-\lambda)\,z'$ corresponds to Minkowski mixing of menus.[54] For any $c \in M$, let Z_c^f be the set of finite menus such that every option $p \in z$ is degenerate and yields consumption c for sure today (i.e., $p = \delta_{(c,w)}$ for some $w \in Z$). For any $z \in Z_c^f$, define $\lambda z \oplus (1-\lambda)\,\delta_{(c',z')} := \left\{\lambda p \oplus (1-\lambda)\,\delta_{(c',z')} : p \in z\right\}$, which is the Minkowski version of ex-post mixing.

Consider a lottery p in a menu $z \in Z_c^f$. Lets mix p and z with a pair (c', z') ex post and call them q and y, respectively (i.e., $q = \lambda p \oplus (1-\lambda)\,\delta_{(c',z')}$ and $y = \lambda z \oplus (1-\lambda)\,\delta_{(c',z')}$). Then $y \in q$ and the independence axiom with respect to the ex-post mixing would state that $\rho_z(p) = \rho_y(q)$. The axiom below strengthens this to independence even with respect to mixtures between z and y.

Axiom 2 (Continuation Linearity)**.** If $p \in z \in Z_c^f$, $y = \lambda z \oplus (1-\lambda)\,\delta_{(c',z')}$ and $q = \lambda p \oplus (1-\lambda)\,\delta_{(c',z')}$ for $c, c' \in M$, $z' \in Z$ and $\lambda \in (0,1)$, then $\rho_z(p) = \rho_{\alpha z+(1-\alpha)y}(\alpha p + (1-\alpha)\,q)$.

The next two axioms are conditions with respect to the classic stationarity axiom originally proposed by Koopmans (1960). In classic stationarity, an agent's choices remain unchanged if all consumptions are delayed by the same number of time periods. Given stochastic

[52] When the outcome space is infinite-dimensional, allowing for all possible vNM utilities would be too permissive and not deliver unique identification.
[53] Note that if the condition is satisfied for N, then it must also be satisfied for all $N' \geq N$ so testing the axiom involves finding a large enough N such that the condition holds.
[54] One could only impose mixing in menus in cases where tomorrow's consumption is the same. The same characterization would then lead to a random utility model where the transition probabilities P_s could also depend on the consumption each period and they all share the same stationary distribution. This could accommodate consumption-dependent stochastic preferences such as habit formation or experimentation.

preferences, classic stationarity would obviously be violated. One way to extend stationarity to a stochastic setup is to require an agent's choice frequencies to remain unchanged if all consumptions are delayed by the same number of time periods: for any $z, y \in Z^f$ and $c \in M$,

$$\rho(z, y) = \rho\left(\delta_{(c,z)}, \delta_{(c,y)}\right).$$

Classic stationarity is normatively appealing and this stochastic stationarity retains much of the flavor of classic stationarity but allows for stochastic choice due to stochastic utilities.[55] However, stochastic stationarity would be violated in our model of ergodic utility. For example, consider the standard utility process, in which the state follows an i.i.d. process (i.e., $P_s = \pi$ for all $s \in S$). Then the agent's choice between the current options could be stochastic. On the other hand, his choice between the delayed options depends on the agent's *expectation* of the discount rate, which is deterministic in this i.i.d. example.[56] Thus, stochastic stationarity will be violated.[57]

Given this observation, we consider two relaxations of Stochastic Stationarity. The first condition, Deterministic Stationarity (Axiom 3) is exactly the classic deterministic stationarity axiom of Koopmans (1960) extended to menus.[58] It states that choices should satisfy stationarity whenever they are deterministic.

Axiom 3 (Deterministic Stationarity)**.** If $\rho(z, y) = 1$, then $\rho\left(\delta_{(c,z)}, \delta_{(c,y)}\right) = 1$ for any $z, y \in Z^f$ and $c \in M$.

The second condition, Average Stationarity (Axiom 4), ensures that stationarity should be satisfied "on average". First, for $\alpha \in [0, 1]$, define the following lottery that yields either the best or worst prize: $p_\alpha := \alpha\delta_{\bar{x}} + (1 - \alpha)\delta_{\underline{x}}$. Thus, p_α is the worst option when $\alpha = 0$ and the best option when $\alpha = 1$. Now, for any $\alpha \in [0, 1]$, one can interpret $\rho(z, p_\alpha)$ as the demand for z relative to p_α, where p_α is the outside option. We can thus interpret

[55] See Lu and Saito (2018) for a stochastic version of the stationarity axiom in a different setup.

[56] Note also that in the i.i.d. example, the discount factor for consumption at period t is given by

$$\beta_{s_1}\mathbb{E}\left[\beta_{s_2}\mathbb{E}\left[\cdots\beta_{s_{t-1}}\mathbb{E}\left[\beta_{s_t}\right]\right]\right] = \beta_{s_1}\mathbb{E}\left[\beta\right]^{t-1} = \beta_{s_1}\delta^{t-1},$$

where $\delta := \mathbb{E}[\beta]$. Interestingly, this corresponds to a model of random quasi-hyperbolic discounting where present bias occurs if $\beta_{s_1} < \delta$ and future bias occurs if $\beta_{s_1} > \delta$.

[57] To see this formally, let p correspond to the option of consuming c_1 today and 0 tomorrow and q correspond to the option of consuming 0 today and c_2 tomorrow. Thus, $\rho(p, q) = \pi\{w(c_1) \geq \beta_{s_1} w(c_2)\}$, which is stochastic and depends on the realization of the stochastic discount rate β_{s_1}. On the other hand, if all consumption is delayed by one period, then $\rho\left(\delta_{(c,p)}, \delta_{(c,q)}\right) = \pi\left\{\cancel{\beta_{s_1}} w(c_1) \geq \cancel{\beta_{s_1}}\delta w(c_2)\right\} = \pi\{w(c_1) \geq \delta w(c_2)\}$, which is not stochastic as $\delta = \mathbb{E}[\beta_{s_2}]$ is deterministic. In general, when realizations and expectations are different, stochastic stationarity will be violated.

[58] It is very similar to the menu stationarity axiom of Higashi et al. (2009) except we only require implication in one direction.

$$\bar{z} := \int_0^1 \rho(z, p_\alpha)\, d\alpha \tag{5}$$

as the "average" demand for z. Notice that this formulation of average demand is similar to the way of measuring consumer surplus by integrating the demand function with respect to price.[59] Average Stationarity says that average demand remains unchanged if all consumptions are delayed by one period.

Axiom 4 (Average Stationarity)**.** For any $z \in Z^f$ and $c \in M$,

$$\int_0^1 \rho(z, p_\alpha)\, d\alpha = \int_0^1 \rho\left(\delta_{(c,z)}, \delta_{(c,p_\alpha)}\right) d\alpha. \tag{6}$$

The axiom can be interpreted as the stationarity on the *surplus* of menus. To see the interpretation, recall from McFadden (1978, 1981) that the surplus of a menu z is given by $\int_S \max_{p\in z} u_s(p)\, d\pi$. Following this definition, the surplus of a menu z delayed by one period is given by $\int_S \left(\int_S \max_{p\in z} u_{s'}(p)\, dP_s\right) d\pi$. If the Markov process is stationary (i.e., $\pi = \int_S P_s d\pi$), then these two surpluses must be the same. This is exactly the implication of the axiom.[60] Thus, Average Stationarity means that the surplus of the menu does not change by the delay.

While Average Stationarity ensures stationarity of the utility process, it does not guarantee ergodicity of the utility process which is crucial for our representation. This is obtained by a final axiom called D-continuity (Axiom 5). First, note that by Monotonicity, if $z \supset y$, then clearly $\rho(z, y) = 1$. By Deterministic Stationarity, this implies that $\rho\left(\delta_{(c,z)}, \delta_{(c,y)}\right) = 1$, which demonstrates classic preference for flexibility. We now require preference for flexibility to be "robust" in the following sense. For any menu $z \in Z$, let $p_{\bar{z}} := \bar{z}\delta_{\bar{x}} + (1-\bar{z})\,\delta_{\underline{x}}$ denote its *probability-equivalent* where $\bar{z}$ is its average demand from equation (5). Since average demand is equivalent to the surplus of the menu, the agent is ex-ante indifferent between the menu and its probability-equivalent. The last axiom states that preference for flexibility is robust even if we perturb the menus z and y slightly by mixing them with the probability-equivalents $p_{\bar{y}}$ and $p_{\bar{z}}$ respectively.

Axiom 5 (D-continuity)**.** There exists $\varepsilon > 0$ such that for any $z, y \in Z$ and $c \in M$, if $z \supset y$ then $\rho\left(\delta_{(c,(1-\varepsilon)z+\varepsilon p_{\bar{y}})}, \delta_{(c,(1-\varepsilon)y+\varepsilon p_{\bar{z}})}\right) = 1$.

---

[59] This is similar to the use of test functions in Lu (2016)

[60] By standard integration by parts, it can be shown that the left hand side of (6) equals to $\int_S \max_{p\in z} u_s(p)\, d\pi$ and the right hand side of (6) equals to $\int_S \left(\int_S \max_{p\in z} u_{s'}(p)\, dP_s\right) d\pi$. This is similar to the use of test functions in Lu (2016).

D-continuity implies that the utility process satisfies Doeblin's condition and is thus ergodic. We are now ready to state our main representation theorem.

Theorem 4. ρ satisfies Axioms 1–5 if and only if it is ergodic.

See Appendix E for the proof. We now provide an outline for the proof of Theorem 4. The first step is the construction of a random expected utility representation where the probability measure is countably additive and continuation menus are evaluated according to the additive linear utility function of Dekel et al. (2001). For this step, we extend Dekel et al. (2001) and Gul and Pesendorfer (2006) with countably-additive probability measures to an infinite-dimensional space. (See Theorem 5 and 6 in Appendix D for the extensions.) Both extensions are known challenges in the literature as the set of utilities over an infinite-dimensional space (without any restrictions) can be no longer compact.[61]

We employ a unified methodology that achieves both. We first focus on finite-dimensional settings. Then, we apply Kolmogorov's extension theorem followed by Tietze extension theorem (Theorem 4.16 of Folland (2013)) by focusing on the set of Lipschitz continuous utilities with common bound;[62] this forms a nice compact set according to the Arzela-Ascoli theorem (see Appendix A). Note that this is not only important for the representation but also crucial for identification in both settings (Theorem 1). In fact, without such a restriction on the set of utilities, identification would not be possible.

Once we have a random expected utility representation where continuation menus are evaluated according to the additive linear functional form, the next step is to show that the random utilities are derived from the stationary distribution of an ergodic utility process. This is where the last three axioms come into play. First, by using Deterministic and Average Stationarity, we show that the random utility is recursive. This allows us to construct a Markov utility process with a stationary distribution that coincides exactly with the distribution of the random utility from the representation. Next, D-continuity ensures that this Markov utility process is ergodic. Finally, the representation is obtained by an application of the Birkhoff ergodic theorem.

Appendix A. Lipschitz Continuous Utilities

Since M and Z are compact metric spaces, $X = M \times Z$ is a compact metric space. Let $C(X)$ denote the set of continuous functions defined on X, $L(X)$ denote the set of Lipschitz continuous functions defined on X, and $L_N(X)$ the set of Lipschitz functions defined on X

[61] For instance, the unit ball is compact in finite-dimensional space but not in infinite-dimensional space. See the discussion after Theorem 3 in Krishna and Sadowski (2014) for more details.

[62] This set is obtained by using the L-continuity (Axiom 1.6).

with Lipschitz bound N. We endow $C(X)$ with the topology of uniform convergence. Fix $\overline{x}, \underline{x} \in X$ and define

$$U_N := \{u \in L_N(X) : 0 = u(\underline{x}) \leq u(x) \leq u(\overline{x}) = 1 \text{ for all } x \in X\}. \tag{7}$$

For each $u \in C(X)$ and $p \in \Delta X$, let $u(p) = \int_X u dp$ denote its expectation. The following result shows that the set of utilities we consider is compact. It is crucial for both characterization and identification, and highlights the role of Lipschitz functions.

Lemma 3. U_N is compact in $C(X)$.

Proof. We will show this using the Arzela-Ascoli Theorem (Theorem 4.43 of Folland (2013)). First, we show that $L_N(X)$ is equicontinuous. Fix $x \in X$ and $\varepsilon > 0$ and consider $y \in X$ such that $|x-y| < \frac{1}{N}\varepsilon$. Thus, for all $u \in L_N(X)$, $|u(x) - u(y)| \leq N|x-y| < \varepsilon$. Since this holds for all $x \in X$, U_N is equicontinuous. Since $0 \leq |u| \leq 1$ for all $u \in U_N$, U_N is pointwise bounded. Next, we show that U_N is closed. Consider $u_k \in U_N$ such that $u_k \to u$. We will show that $u \in U_N$. Since u_k is bounded, we have $u(x) - u(y) = \lim_k (u_k(x) - u_k(y)) \leq \lim_k N|x-y| = N|x-y|$ for all $x, y \in X$. Thus, $u \in L_N(X)$. Next, note that for all k,$0 = u_k(\underline{x}) \leq u_k(x) \leq u_k(\bar{x}) = 1$, so $0 = u(\underline{x}) \leq u(x) \leq u(\bar{x}) = 1$.This shows $u \in U_N$, hence U_N is closed. By the Arzela-Ascoli (Theorem 4.43 of Folland (2013)) , U_N is compact in $C(X)$. ■

A.1. **Proof of Lemma 1.** We first show the following lemma which characterizes distributions on a compact subset of $C(X)$.

Lemma 4. Let $\mu, \nu \in \Delta U$ where U is a compact subset of $C(X)$. If $\int_U e^{ru(p)} d\mu = \int_U e^{ru(p)} d\nu$ for all $r \geq 0$ and $p \in \Delta X$, then $\mu = \nu$.

Proof. Let Φ denote the set of continuous functions ϕ defined on U such that $\phi(u) = \sum_{i=1}^n a_i e^{r_i u(p_i)}$ for some n, $a_i \in \mathbb{R}$, $r_i \geq 0$ and $p_i \in \Delta X$ for each $i \in \{1, \ldots, n\}$. Thus, for all $\phi \in \Phi$, $\int_U \phi(u) d\mu = \int_U \sum_{i=1}^n a_i e^{r_i u(p_i)} d\mu = \int_U \sum_{i=1}^n a_i e^{r_i u(p_i)} d\nu = \int_U \phi(u) d\nu$. We will show that Φ is uniformly dense in $C(U)$ by the Stone-Weierstrass Theorem (Theorem 9.13 of Aliprantis and Border (2006) (henceforth, AB)). First note that Φ is a vector space that includes constants since $e^{0u(p)} = 1 \in \Phi$. It is easy to show that (i) Φ is closed under multiplication and (ii) Φ separates points in U.[63]

Since U is compact, Φ is a subalgebra, contains the constant function and separates points in U, Φ is uniformly dense in $C(U)$ by the Stone-Weierstrass Theorem. This means that for

[63] To show (i), consider $a_1 e^{r_1 u(p_1)}, a_2 e^{r_1 u(p_2)} \in \Phi$. If $r_1 + r_2 > 0$, then $a_1 e^{r_1 u(p_1)} a_2 e^{r_2 u(p_2)} = a_1 a_2 e^{(r_1+r_2) u\left(\frac{r_1}{r_1+r_2} p_1 + \left(1 - \frac{r_1}{r_1+r_2}\right) p_2\right)} \in \Phi$. On the other hand, if $r_1 + r_2 = 0$, then $r_1 = r_2 = 0$ and $a_1 e^{r_1 u(p_1)} a_2 e^{r_2 u(p_2)} = a_1 a_2 \in \Phi$. Thus, Φ is closed under multiplication. To show (ii), suppose $u, v \in U$ such that $u \neq v$. Thus, there is some $x \in X$ such that $u(x) > v(x)$ without loss of generality. If we let $p = \delta_x$, then $u(p) = u(x) > v(x) = v(p)$ so $e^{u(p)} > e^{v(p)}$. Thus, Φ separates points in U.

any $\phi \in C(U)$, we can find $\phi_k \in \Phi$ such that $\phi_k \to \phi$ uniformly. Hence, if we fix some $\varepsilon > 0$, then there exists some n such that $|\phi_k - \phi| \leq \varepsilon$ for all $k > n$. This implies that for all $u \in U$, $\phi_k(u) \leq |\phi_k(u) - \phi(u)| + |\phi(u)| \leq |\phi(u)| + \varepsilon$. Thus, ϕ_k are all dominated by a integrable function, so by dominated convergence, $\int_U \phi(u)\, d\mu = \lim_k \int_U \phi_k(u)\, d\mu = \lim_k \int_U \phi_k(u)\, d\nu = \int_U \phi(u)\, d\nu$. By AB Theorem 15.1, $\mu = \nu$. ■

We now prove Lemma 1. Define the mapping $\xi : S \to U$ as in equation (3), or $\xi_s(c, z) = \phi_s\left(c, \int_S \sup_{p \in z} u_{\tilde{s}}(p)\, dP_s\right)$. Consider two states $s, s' \in S$ such that $\xi_s = \xi_{s'}$. We will show that this means that $P_s \circ \xi^{-1} = P_{s'} \circ \xi^{-1}$. Let $\nu = P_s \circ \xi^{-1}$, $\nu' = P_{s'} \circ \xi^{-1}$ and $z = \left\{p, \alpha\delta_{\bar{x}} + (1-\alpha)\delta_{\underline{x}}\right\}$. Since $\xi_s = \xi_{s'}$ and $\phi_\cdot$ is strictly increasing in the second argument, we have $\int_U \max\{u(p), \alpha\}\, d\nu = \int_S \sup_{p \in z} u_{\tilde{s}}(p)\, dP_s = \int_S \sup_{p \in z} u_{\tilde{s}}(p)\, dP_{s'} = \int_U \max\{u(p), \alpha\}\, d\nu'$ for any $\alpha \in [0,1]$. By Theorem 1.57 of Müller and Stoyan (2002), for any increasing convex function φ, $\int_U \varphi(u(p))\, d\nu = \int_U \varphi(u(p))\, d\nu'$. Thus by Lemma 4, $\nu = \nu'$ because ν and ν' are probability measures on U_N, which is compact by Lemma 3. We can now define a transition kernel ν_v on U such that $\nu_v := P_s \circ \xi^{-1}$ where $v = u_s$. If we let $\mu = \pi \circ \xi^{-1}$, then $\int_U \nu_v(B)\, d\mu = \int_S \nu_{u_s}(B)\, d\pi = \int_S P_s(\xi^{-1}(B))\, d\pi = \pi(\xi^{-1}(B)) = \mu(B)$, where the first and the last equality hold by the definition of μ, the second equality holds by definition of ν_v, and the third equality holds because π is a stationary distribution of P. Thus, the utility process is a stationary Markov process. Moreover, for any measurable B, we have μ-a.s. $\nu_v(B) = P_s(\xi^{-1}(B)) \geq \delta\pi(\xi^{-1}(B)) = \delta\mu(B)$ so the Markov process satisfies Doeblin's condition and is thus ergodic.

Appendix B. Proof of Theorem 1 (Uniqueness)

From Lemma 1, the utility process is ergodic so let μ and μ' denote the stationary utility distributions for ρ and ρ' respectively. For every $z = \{p, q\} \in Z^*$, we have $\rho_z(p) = \lim_{n\to\infty} \frac{1}{n} \sum_{k \geq 0} 1_{B(p,z)}(s_{tk+1}) = \mu\{u \in U : u(p) \geq u(q)\}$ and likewise for ρ' and μ', where the first equality is by the ergodic representation and the second equality is by the Birkoff ergodic theorem. Choose any binary menu $z = \{p, q\} \in Z$. For each $t \in T$, define $p^t = p_{z^t,t}$ and $q^t = q_{z^t,t}$. Then $p^t \to p$ and $q^t \to q$ as $t \to \infty$. By definition $z^t = \{p^t, q^t\} \in Z^*$ and $z^t \to z$ by Lemma 2.

Step 1: If $u(p) = u(q)$ with μ-measure zero, then $\lim_{t\to\infty} \rho(p^t, q^t) = \mu\{u(p) \geq u(q)\}$.

Proof. First, note that μ-a.s. $\lim_t 1_{u(p^t) \geq u(q^t)} = 1_{u(p) \geq u(q)}$.[64] By the dominated convergence theorem, we thus have $\lim_t \rho(p^t, q^t) = \lim_t \int_U 1_{u(p^t) \geq u(q^t)} d\mu = \int_U 1_{u(p) \geq u(q)} d\mu = \mu\{u(p) \geq u(q)\}$, as desired. □

[64] To see why, first suppose $u(p) \geq u(q)$, but $\liminf_t 1_{u(p^t) \geq u(q^t)} = 0$. Thus, we can find a subsequence p^k, q^k such that $u(p^k) < u(q^k)$ so $u(p) \leq u(q)$ yielding a contradiction as $u(p) \neq u(q)$ μ-a.s.. On the other hand, if $u(p) < u(q)$, then clearly $\limsup_t 1_{u(p^t) \geq u(q^t)} = 0$.

Step 2: If $u(p) = u(q)$ with μ'-a.s., then $u(p) = u(q)$ with μ-a.s.

Proof. Let $q = p_\alpha := \alpha\delta_{\bar{x}} + (1-\alpha)\delta_{\underline{x}}$ and suppose that $u(p) = u(q) = \alpha$ μ'-a.s. We will show that this implies $u(p) = \alpha$ μ-a.s. Fix a positive number ε. Consider $p_{\alpha+\varepsilon}$ and $p_{\alpha-\varepsilon}$ and note that $u(p_{\alpha+\varepsilon}) > u(p) > u(p_{\alpha-\varepsilon})$ μ'-a.s. for all $\varepsilon > 0$. By regularity, without loss of generality, we can choose ε such that $u(p) = u(p_{\alpha+\varepsilon})$ and $u(p) = u(p_{\alpha-\varepsilon})$ with μ-measure zero. Thus, $\mu\{u(p) \geq u(p_{\alpha-\varepsilon})\} = \lim_t \rho\left(p^t, p^t_{\alpha-\varepsilon}\right) = \lim_t \rho'\left(p^t, p^t_{\alpha-\varepsilon}\right) = \mu'\{u(p) \geq u(p_{\alpha-\varepsilon})\} = 1$, where the first and the third equality hold by Step 1, the second equality holds by the supposition of Theorem 1 that ρ and ρ' coincide on binary sets, and the last equality holds by the supposition that $u(p) = \alpha$ μ'-a.s.. By the symmetric argument for p and $p_{\alpha+\varepsilon}$, $\mu\{u(p_{\alpha+\varepsilon}) \geq u(p)\} = \lim_t \rho\left(p^t_{\alpha+\varepsilon}, p^t\right) = \lim_t \rho'\left(p^t_{\alpha+\varepsilon}, p^t\right) = \mu'\{u(p_{\alpha+\varepsilon}) \geq u(p)\} = 1$. Thus, $u(p) \in [\alpha-\varepsilon, \alpha+\varepsilon]$ μ-a.s. Since ε is an arbitrary positive number, $u(p) = \alpha$ μ-a.s. as desired. □

Step 3: For any $p \in \Delta(X)$, $u(p)$ has the same distribution under μ and under μ'.

Proof. Fix any $p \in \Delta(X)$ and $\alpha \in \mathbb{R}$ to show $\mu\{u(p) \geq \alpha\} = \mu'\{u(p) \geq \alpha\}$. By the regularity of μ, it suffices to consider the following two cases.

Case 1: The case when $\mu\{u(p) = \alpha\} = 0$. Let $q = p_\alpha := \alpha\delta_{\bar{x}} + (1-\alpha)\delta_{\underline{x}}$. By Step 1, $\mu\{u(p) \geq \alpha\} = \lim_{t\to\infty} \rho(\rho^t, q) = \lim_{t\to\infty} \rho'(\rho^t, q) = \mu'\{u(p) \geq \alpha\}$.

Case 2: The case when $\mu\{u(p) = \alpha\} = 1$. By Step 2, $\mu'\{u(p) = \alpha\} = 1 = \mu\{u(p) = \alpha\}$. □

Now, by Step 3, $\int_U e^{ru(p)} d\mu = \int_U e^{ru(p)} d\mu'$ for all $r \geq 0$ and $p \in \Delta X$. Since μ and μ' are probability measure on U_N, which is compact by Lemma 3. Thus, $\mu = \mu'$ by Lemma 4. Since each $u \in U$ determines the transition kernel on U, this means that the Markov utility process induced by μ and μ' are the same. The converse is trivial.

Appendix C. Proof of Theorem 2 and 3 (Characterization of ICM)

C.1. **Proof of Theorem 2.** Let μ denote the stationarity distribution for the utility process so by the ergodic theorem, we have $\rho(p, q) = \mu\{u \in U : u(p) \geq u(q)\}$ for all $\{p, r\} \in Z^*$. That additive separability implies 1-ICM and the standard utility implies both 1-ICM and 2-ICM are straightforward. We now show that 1-ICM implies the utility process is additively separable. Fix $z, y \in Z$. Consider $\{p, r\} \in Z^*$ such that $p = \frac{1}{4}\delta_{(0,z)} + \frac{1}{4}\delta_{(0,y)} + \frac{1}{4}\delta_{(c,z)} + \frac{1}{4}\delta_{(c,y)}, r = \frac{1}{2}\delta_{(0,z)} + \frac{1}{2}\delta_{(c,y)}$. Note that $p_M = r_M = \frac{1}{2}\delta_0 + \frac{1}{2}\delta_c$ and $p_Z = r_Z = \frac{1}{2}\delta_z + \frac{1}{2}\delta_y$. Let $\{q\} \in Z^*$ denote the singleton 1-period menu such that $q_M = \frac{1}{2}\delta_0 + \frac{1}{2}\delta_c = p_M$ and $q_Z = \frac{1}{2}\delta_z + \frac{1}{2}\delta_y$. By 1-ICM, we thus have $1 = \rho_{\{q\}}(q) = \rho(p, r) = \rho(r, p)$. Thus a.s. $\frac{1}{4}u(0, z) + \frac{1}{4}u(0, y) + \frac{1}{4}u(c, z) + \frac{1}{4}u(c, y) = \frac{1}{2}u(0, z) + \frac{1}{2}u(c, y)$. That is, $u(0, y) + u(c, z) = u(0, z) + u(c, y)$.

Let $y = \underline{x}^t \to \underline{x}$ so $u(0, \underline{x}^t) \to u(\underline{x}) = 0$. If we let $v_s(z) := \mathbb{E}_s\left[\sup_{p \in z} u_{s'}(p)\right]$ then $v_s(\underline{x}^t) \to 0$. Thus, we have a.s. $\phi_s(c, v_s(z)) = \phi_s(0, v_s(z)) + \phi_s(c, 0)$. Letting $w_s(c) := \phi_s(c, 0)$ and $\beta_s(v) = \phi_s(0, v)$, we have a.s. $\phi_s(c, v) = w_s(c) + \beta_s(v)$ as desired.

We now show that imposing 2-ICM in addition to 1-ICM implies the utility process must be standard. By 1-ICM, we have a.s. $\phi_s(c, v) = w_s(c) + \beta_s(v)$ from above. Consider $\{p, r\} \in Z^*$ such that $p = \frac{1}{4}\delta_{(0,p_0)} + \frac{1}{4}\delta_{\left(0,\delta_{(0,y)}\right)} + \frac{1}{4}\delta_{(0,q_0)} + \frac{1}{4}\delta_{(0,r_0)}, r = \frac{1}{2}\delta_{(0,p_0)} + \frac{1}{2}\delta_{(0,r_0)}$, where $p_0 = b\delta_{(0,z)} + (1-b)\,\delta_{(0,y)}$, $q_0 = ab\delta_{(m,z)} + a(1-b)\,\delta_{(m,y)} + (1-a)\,b\delta_{(0,z)} + (1-a)(1-b)\,\delta_{(0,y)}$, and $r_0 = a\delta_{(m,y)} + (1-a)\,\delta_{(0,y)}$ for $a, b \in [0,1]$. Note that the distribution of 2-period consumptions of p and r are $\frac{1}{2}\delta_{(0,\delta_0)} + \frac{1}{2}\delta_{(0,a\delta_m+(1-a)\delta_0)}$ while their menu distributions are $\frac{1}{2}\delta_{(b\delta_z+(1-b)\delta_y)} + \frac{1}{2}\delta_{\delta_y}$. Thus, by 2-ICM, $1 = \rho(p,r) = \rho(r,p)$. Hence, we have a.s. $\frac{1}{4}u(0,p_0) + \frac{1}{4}u\left(0, \delta_{(0,y)}\right) + \frac{1}{4}u(0,q_0) + \frac{1}{4}u(0,r_0) = \frac{1}{2}u(0,p_0) + \frac{1}{2}u(0,r_0)$. That is, $u\left(0, \delta_{(0,y)}\right) + u(0,q_0) = u(0,p_0) + u(0,r_0)$. Thus, we have a.s. $\beta_s(\mathbb{E}_s[u_{s'}(0,y)]) + \beta_s(\mathbb{E}_s[u_{s'}(q_0)]) = \beta_s(\mathbb{E}_s[u_{s'}(p_0)]) + \beta_s(\mathbb{E}_s[u_{s'}(r_0)])$. Let $y = \{\underline{x}^t\} \to \{\underline{x}\}$ and $z = \bar{x}^t \to \bar{x}$ so $u_s(0, \underline{x}^t) \to u_s(\underline{x}) = 0$ and $v_s(\bar{x}^t) \to 1$. By definition, $\beta_s(0) = \phi_s(0,0) = 0$ and $w_s(0) = \phi_s(0,0) = 0$. Thus we have a.s. $\beta_s(\mathbb{E}_s[aw_{s'}(m) + b\beta_{s'}(1)]) = \beta_s(\mathbb{E}_s[aw_{s'}(m)]) + \beta_s(\mathbb{E}_s[b\beta_{s'}(1)])$, or $\beta_s(a\mathbb{E}_s[w_{s'}(m)] + b\mathbb{E}_s[\beta_{s'}(1)]) = \beta_s(a\mathbb{E}_s[w_{s'}(m)]) + \beta_s(b\mathbb{E}_s[\beta_{s'}(1)])$ for all $a, b \in [0,1]$. Let $\xi_s := \min\{\mathbb{E}_s[w_{s'}(m)], \mathbb{E}_s[\beta_{s'}(1)]\}$. Since $\mathbb{E}_s[w_{s'}(m)] + \mathbb{E}_s[\beta_{s'}(1)] = \mathbb{E}_s[w_{s'}(m) + \beta_{s'}(1)] = \mathbb{E}_s[\phi_s(m,1)] = 1$, it follows that $\xi_s > 0$. From the equation above, we have $\beta_s(x+y) = \beta_s(x) + \beta_s(y)$ for all $x, y \in [0, \xi_s]$. This is a Cauchy functional equation with bounded domain, and since β_s is continuous, we have a.s. $\beta_s(x) = \beta_s x$ for all $x \in [0, \xi_s]$ where β_s is a constant (see pg. 45 of Aczel (1966)). Now, for $v \in [0, 2\xi_s]$, $\beta_s(v) = \beta_s\left(\frac{v}{2} + \frac{v}{2}\right) = 2\beta_s\left(\frac{v}{2}\right) = \beta_s v$. By iteration, we have $\beta_s(v) = \beta_s v$ for all $v \in [0,1]$ as desired.

C.2. **Definition of Repeated Independence (RI).** In the main part of the paper, we explained how to mix 1-period menus with a lottery $r \in \Delta M$. In this subsection, we formally define how to mix simple t-period menus. Fix some t-period menu $z \in Z^*$ that is also t-simple. For any lottery p that yields z in $t' \le t$ periods, we will define $r_{t'}(p)$ as the t'-times repeated mixture between p and $r \in \Delta M$. This is constructed as follows. First, define $r_1(\cdot)$ exactly as in the 1-period case where for every p, $r_1(p) = (\alpha p_M + (1-\alpha)\,r; \alpha z \otimes (1-\alpha)\,r)$, where p_M is the marginal distribution of p on M. Now for $1 < t' \le t$, we will recursively define $r_{t'}(\cdot)$. First, given $r_{t'-1}(\cdot)$ and some p, define two lotteries $\hat{p}$ and $\hat{r}$ such that $\hat{p}(A \times B) := p\left(A \times r_{t'-1}^{-1}(B)\right)$ and $\hat{r}(A \times B) := r(A)\,p\left(\Delta M \times r_{t'-1}^{-1}(B)\right)$ for all measurable A and B. Note that $\hat{p}$ and $\hat{r}$ are the continuation lotteries where all future lotteries are also mixed with r. Next, define $r_{t'}(p) = \alpha\hat{p} + (1-\alpha)\,\hat{r}$. Finally, set $\alpha p \otimes (1-\alpha)\,r = r_t(p)$ and $\alpha z \otimes (1-\alpha)\,r := \{\alpha p \otimes (1-\alpha)\,r : p \in z\}$.

C.3. **Proof of Theorem 3.** Note that by Theorem 2, all we need to show is that the utility process is standard if and only if ρ satisfies IRU and RI. Since the standard utility process trivially satisfies IRU and RI, we will show the converse. By IRU, we have $\phi_s(c,v) = w_s(c) + \beta_s(c)v$. Note that $0 = \phi_s(0,0) = w_s(0)$ and $1 = \phi_s(m,1) = w_s(m) + \beta_s(m)$. Consider a 2-period $z = \{p_0, q_0\} \in Z^*$ where $p_0 = \frac{1}{2}\delta_{(c_1,p)} + \frac{1}{2}\delta_{\left(c_2,\delta_{(c_2,z)}\right)}$,$q_0 = \frac{1}{2}\delta_{(c_1,q)} + \frac{1}{2}\delta_{(c_2,q)}$, $p = \lambda_1\delta_{(m,z)} + (1-\lambda_1)\delta_{(c_2,z)}$ and $q = \lambda_2\delta_{(m,z)} + (1-\lambda_2)\delta_{(c_2,z)}$ for $c_1, c_2 \in (0,m)$. Note also that $u_s(p_0) = w_s\left(\frac{1}{2}\delta_{c_1} + \frac{1}{2}\delta_{c_2}\right) + \frac{1}{2}\beta_s(c_1)\mathbb{E}_s[u_{s'}(p)] + \frac{1}{2}\beta_s(c_2)\mathbb{E}_s[u_{s'}(c_2,z)]$ and $u_s(q_0) = w_s\left(\frac{1}{2}\delta_{c_1} + \frac{1}{2}\delta_{c_2}\right) + \left(\frac{1}{2}\beta_s(c_1) + \frac{1}{2}\beta_s(c_2)\right)\mathbb{E}_s[u_{s'}(q)]$. To simplify notation, let $\beta_i := \beta_s(c_i)$ and $\tilde{u}_s := \mathbb{E}_s[u_{s'}]$. Now, by definitions $u_s(p_0) \geq u_s(q_0) \Leftrightarrow \beta_1\lambda_1 \geq (\beta_1+\beta_2)\lambda_2$.[65]

Let $r = \delta_{c_2}$ and consider the 2-period $z' = az \otimes (1-a)r \in Z^*$. Note that $z' = \{ap_0 \otimes (1-a)r, aq_0 \otimes (1-a)r\}$ where $ap_0 \otimes (1-a)r = \frac{1}{2}\left(a\delta_{(c_1,p')} + (1-a)\delta_{(c_2,p')}\right) + \frac{1}{2}\delta_{\left(c_2,\delta_{(c_2,z')}\right)}$, $aq_0 \otimes (1-a)r = \frac{1}{2}\left(a\delta_{(c_1,q')} + (1-a)\delta_{(c_2,q')}\right) + \frac{1}{2}\delta_{(c_2,q')}$, $p' = a\lambda_1\delta_{(m,z')} + (1-a\lambda_1)\delta_{(c_2,z')}$, and $q' = a\lambda_2\delta_{(m,z')} + (1-a\lambda_2)\delta_{(c_2,z')}$. Note also that $u_s(ap_0 \otimes (1-a)r) = w_s\left(\frac{a}{2}\delta_{c_1} + \left(1-\frac{a}{2}\right)\delta_{c_2}\right) + \frac{1}{2}(a\beta_1 + (1-a)\beta_2)\tilde{u}_s(p') + \frac{1}{2}\beta_2\tilde{u}_s(c_2,z')$ and $u_s(aq_0 \otimes (1-a)r) = w_s\left(\frac{a}{2}\delta_{c_1} + \left(1-\frac{a}{2}\right)\delta_{c_2}\right) + \left(\frac{a}{2}\beta_1 + \left(1-\frac{a}{2}\right)\beta_2\right)\tilde{u}_s(q')$. For simplicity, let $\beta_a := a\beta_1 + (1-a)\beta_2$ and recall $\tilde{u}_s = \mathbb{E}_s[u_{s'}]$. Now by definitions, $u_s(ap_0 \otimes (1-a)r) \geq u_s(aq_0 \otimes (1-a)r) \Leftrightarrow \beta_a\lambda_1 \geq (\beta_a+\beta_2)\lambda_2$.[66]

By RI, $\mu\left\{\frac{\beta_a}{\beta_a+\beta_2} \geq \frac{\lambda_2}{\lambda_1}\right\} = \rho(ap_0 \otimes (1-a)r, aq_0 \otimes (1-a)r) = \rho(p_0,q_0) = \mu\left\{\frac{\beta_1}{\beta_1+\beta_2} \geq \frac{\lambda_2}{\lambda_1}\right\}$. Since this is true for all $\lambda_1, \lambda_2 \in (0,1)$, it must be that $\frac{\beta_1}{\beta_1+\beta_2}$ and $\frac{\beta_a}{\beta_a+\beta_2}$ have the same distribution for all $a > 0$. If we let $\xi := \frac{\beta_1}{\beta_2}$, then ξ has the same distribution as $\frac{\beta_a}{\beta_2} = a\xi + (1-a)$. Equivalently, this implies that $\xi - 1$ has the same distribution as $a(\xi-1)$ for all $a > 0$. Let κ be the infimum of the support of $\xi - 1$. Since $\frac{\beta_1}{\beta_2} \geq 0$, $\kappa \geq -1$. Since $\xi - 1$ and $a(\xi-1)$ have the same distribution, it must be that $\kappa = 0$. Thus, a.s. $0 \leq \xi - 1 = \frac{\beta_1}{\beta_2} - 1$, or $\beta_s(c_1) = \beta_1 \geq \beta_2 = \beta_s(c_2)$ a.s. Since this was for arbitrary $c_1, c_2 \in (0,m)$, it must be that $\beta(c_1) = \beta(c_2)$ for all $c_1, c_2 \in (0,m)$. Continuity of β then yields β must be constant on M.

Appendix D. Extensions of DLR and GP

In this section, we present extensions of the main representation theorems of Dekel et al. (2001) (henceforth DLR) and Gul and Pesendorfer (2006) (henceforth GP).[67] Using the same methodology, we extend their results to countably-additive probability measures in

[65] To see this, $u_s(p_0) \geq u_s(q_0) \Leftrightarrow \beta_1\tilde{u}_s(p) + \beta_2\tilde{u}_s(c_2,z) \geq (\beta_1+\beta_2)\tilde{u}_s(q) \Leftrightarrow \beta_1(\tilde{u}_s(p) - \tilde{u}_s(q)) \geq \beta_2(\tilde{u}_s(q) - \tilde{u}_s(c_2,z)) \Leftrightarrow \beta_1(\lambda_1-\lambda_2)(\tilde{u}_s(m,z) - \tilde{u}_s(c_2,z)) \geq \beta_2\lambda_2(\tilde{u}_s(m,z) - \tilde{u}_s(c_2,z)) \Leftrightarrow \beta_1\lambda_1 \geq (\beta_1+\beta_2)\lambda_2$, where the last inequality follows from the fact that $\tilde{u}_s(m,z) \geq \tilde{u}_s(c_2,z)$ a.s. as $m \geq c_2$.

[66] To see this, $u_s(ap_0 \otimes (1-a)r) \geq u_s(aq_0 \otimes (1-a)r) \Leftrightarrow \beta_a\tilde{u}_s(p') + \beta_2\tilde{u}_s(c_2,z') \geq (\beta_a+\beta_2)\tilde{u}_s(q') \Leftrightarrow \beta_a(\tilde{u}_s(p') - \tilde{u}_s(q')) \geq \beta_2(\tilde{u}_s(q') - \tilde{u}_s(c_2,z')) \Leftrightarrow \beta_a a(\lambda_1-\lambda_2)(\tilde{u}_s(m,z') - \tilde{u}_s(c_2,z')) \geq \beta_2 a\lambda_2((\tilde{u}_s(m,z') - \tilde{u}_s(c_2,z'))) \Leftrightarrow \beta_a\lambda_1 \geq (\beta_a+\beta_2)\lambda_2$, where the last inequality follows from the fact that $\tilde{u}_s(m,z) \geq \tilde{u}_s(c_2,z)$ a.s.

[67] See also Dekel et al. (2007).

infinite-dimensional settings. In both cases, we achieve this by using Lipschitz continuous utilities with a common bound. We first focus on finite-dimensional settings and then apply Kolmogorov's extension theorem followed by Tietze extension theorem. On an abstract level, this is analogous to the extension to uniformly continuous paths for the construction of Brownian motion.[68]

In this section, let X be a compact metric space and U_N be the set of Lipschitz continuous utilities with common bound N as defined by (7). We assume that X contains two elements $\bar{x}$ and $\underline{x}$. Let $Z = \mathcal{K}(\Delta X)$ denote the set of non-empty compact subsets of ΔX.

D.1. **Extension of DLR.** Consider a binary relation $\succeq$ on Z and the following conditions[69]: **C1.1** (Weak Order): $\succeq$ is complete and transitive; **C1.2** (Flexibility): $z \subset y$ implies $z \preceq y$ for all $z, y \in Z$; **C1.3** (Independence): $z \succeq y$ implies $\alpha z + (1-\alpha) w \succeq \alpha y + (1-\alpha) w$ for all $\alpha \in (0,1)$ and $z, y, w \in Z$; **C1.4** (Continuity): $\succeq$ is continuous; **C1.5** (Best-Worst): $\bar{x} \succeq \{x, \bar{x}\}$ and $x \succeq \{x, \underline{x}\}$ for all $x \in X$; **C1.6** (L-Continuity): There exists $N > 0$ such that for all $x_1, x_2 \in X$ and $\alpha \in [0,1]$, $|x_1 - x_2| \leq \frac{\alpha}{N}$ implies $(1-\alpha)\delta_{x_1} + \alpha\delta_{\bar{x}} \succeq \left\{(1-\alpha)\delta_{x_1} + \alpha\delta_{\bar{x}}, (1-\alpha)\delta_{x_2} + \alpha\delta_{\underline{x}}\right\}$.

Theorem 5 (DLR extension)**.** $\succeq$ satisfies Condition C1.1–1.6 if and only if there exists a countable additive probability measure ν on U_N such that $\succeq$ is represented by the function $v : Z \to \mathbb{R}$ where $v(z) = \int_{U_N} \sup_{p \in z} u(p)\, d\nu$.

We now prove Theorems 5. The necessity of the axioms is straightforward.[70] We now show sufficiency. Since X is separable, let $X^* \subset X$ be a countable dense subset of X and without loss of generality, assume $\underline{x}, \bar{x} \in X^*$. We first show three lemmas. The following preliminary lemma modified from Dekel et al. (2007) characterizes Lipschitz continuous functions on X^*. See Section S.3 for the proof

Lemma 5. Let X^* be a dense subset of X and suppose $v : X^* \to \mathbb{R}$ is such that $v(\bar{x}) = 1$ and $v(\underline{x}) = 0$. Then the following statements (i) and (ii) are equivalent: (i) There exist $N > 0$ such that if $|x_1 - x_2| \leq \frac{\alpha}{N}$ for $x_1, x_2 \in X^*$ and $\alpha \in [0,1]$, then $\alpha v(\bar{x}) + (1-\alpha) v(x_1) \geq \alpha v(\underline{x}) + (1-\alpha) v(x_2)$; (ii) v is Lipschitz continuous with bound N.

The next lemma is essential. In fact, it shows the result for any finite $W \subset X^*$.

[68] Other papers that also employ Kolmogorov's extension in this manner include Lu and Saito (2018), who do not address the continuity of utilities, and Frick et al. (2018), who obtain a measure with finite support (ignoring ties).
[69] While DLR formally considers all non-empty subsets of ΔX, it is without loss to focus on compact sets.
[70] Condition C1.1–1.4 follow from the same arguments as in DLR. It is easy to see C1.5 from the representation while C1.6 follows from Lemma 5 above.

Lemma 6. There exists a probability measure ν on U_N such that for all finite $W \subset X^*$, the function $v : Z \to \mathbb{R}$ where $v(z) = \int_{U_N} \sup_{p \in z} u(p)\, d\nu$ represents $\succeq$ on $\mathcal{K}(\Delta W)$.

Proof. We prove this in a series of steps.

Step 1: There exists a measure μ on $\mathbb{R}^{X^*}$ such that $\int_{\mathbb{R}^{X^*}} \sup_{p \in z} u(p)\, d\mu$ represents $\succeq$ on $\mathcal{K}(\Delta W)$ for all finite $W \subset X^*$.

Proof. From DLR, Condition C1.1–C1.4 imply that for each finite $W \subset X^*$ where $\underline{x}, \bar{x} \in W$, there exists a probability measure μ_W on $\mathbb{R}^W$ such that $\int_{\mathbb{R}^W} \sup_{p \in z} u(p)\, d\mu_W$ represents $\succeq$ on $\mathcal{K}(\Delta W)$. By C1.5, we have $u(\underline{x}) \leq u(x) \leq u(\bar{x})$ μ_W-a.s. for all $x \in X$. Thus, we can assume that $u(\bar{x}) = 1$ and $u(\underline{x}) = 0$ without loss of generality. With this normalization of utilities, the DLR representation is unique so all these μ_W are consistent. By Kolmogorov's extension, the result holds. □

Step 2: There exists $N > 0$ such that μ-a.s. for all $\alpha \in [0,1]$ and $x_1, x_2 \in X^*$,$|x_1 - x_2| \leq \alpha/N$ implies $\alpha + (1-\alpha)\, u(x_1) \geq (1-\alpha)\, u(x_2)$.

Proof. By C1.6, N satisfying the condition exists. For $\alpha \in [0,1]$ and $x_1, x_2 \in X^*$, define $U_\alpha^{x_1,x_2} := \{u \in \mathbb{R}^{X^*} : |x_1 - x_2| \leq \frac{\alpha}{N} \implies \alpha + (1-\alpha)\, u(x_1) \geq (1-\alpha)\, u(x_2)\}$. By C1.6, $\mu(U_\alpha^{x_1,x_2}) = 1$ for all $\alpha \in [0,1]$ and $x_1, x_2 \in X^*$. Let $U_\alpha := \bigcap_{x_1,x_2 \in X^*} U_\alpha^{x_1,x_2}$ so by the countable additivity of μ and the fact that X^* is a countable dense subset of X, $\mu(U_\alpha) = 1$ for any $\alpha \in [0,1]$. Let I^* be the rationals in $[0,1]$ so by the same argument, $\mu(\bigcap_{\alpha \in I^*} U_\alpha) = 1$. We want to show that $\mu\left(\bigcap_{\alpha \in [0,1)} U_\alpha\right) = 1$. This holds because $\bigcap_{\alpha \in I^*} U_\alpha \subset \bigcap_{\alpha \in [0,1)} U_\alpha$.[71] □

By Step 2, Lemma 5 yields $\mu(L_N(X^*)) = 1$. By the Lipschitz version of the Tietze extension theorem (see McShane (1934)), we can extend μ on $L_N(X^*)$ to a probability measure ν on $L_N(X)$. Note that Lipschitz continuity ensures this is possible, similar to the role of uniform continuity in the construction of Brownian motion. Finally, we show $\nu(U_N) = 1$. For each $x \in X$, define $U_x := \{u \in L_N(X) : 0 = u(\underline{x}) \leq u(x) \leq u(\bar{x}) = 1\}$. It follows from the countable additivity of ν and the separability of X that $\nu(\bigcap_{x \in X^*} U_x) = 1$. Then by the continuity of u, it is easy to show $\cap_{x \in X^*} U_x \subset \cap_{x \in X} U_x$; thus $\nu(U_N) = 1$.[72] For any $z \in Z$, define $v(z) := \int_{U_N} \sup_{p \in z} u(p)\, d\nu$. Thus, v represents $\succeq$ on $\mathcal{K}(\Delta W)$ for all finite $W \subset X^*$. ■

[71] We will show that for any $u \in \bigcap_{\alpha \in I^*} U_\alpha$ and $\alpha \in [0,1)$, $u \in U_\alpha$. Choose any $x_1, x_2 \in X^*$ such that $|x_2 - x_1| \leq \alpha/N$ and consider a sequence α_k of I^* such that $\alpha_k \to \alpha$ and $\alpha_k \geq \alpha$. Since $|x_2 - x_1| \leq \alpha_k/N$ and $u \in \bigcap_{\alpha \in I^*} U_\alpha$, we have $u(x_2) - u(x_1) \leq \alpha_k/(1-\alpha_k)$ for each k. Since $\alpha_k \to \alpha$, we have $u(x_2) - u(x_1) \leq \alpha/(1-\alpha)$ so $u \in U_\alpha$. Thus, $\mu(\bigcap_{\alpha \in [0,1]} U_\alpha) = 1$ as desired.

[72] Suppose that $u \in \cap_{x \in X^*} U_x$ and consider $x \in X$. If $x \in X^*$, then the result holds trivially so suppose $x \notin X^*$. Since X^* is dense in X, there exists a sequence x_k of X^* such that $x_k \to x$. Since $0 \leq u(x_k) \leq 1$ for each k, we have $0 \leq u(x) \leq 1$ by the continuity of u.

Since X^* is dense and Dirac measures are extreme points in ΔX, the next lemma follows from the Krein-Milman theorem (AB Theorem 15.10):

Lemma 7. For every $p \in \Delta X$, there exists a sequence $p_n \to p$ such that each p_n has a finite support in X^*.

We now complete the proof by showing that v represents $\succeq$ on Z. First we show that v is continuous. Note that $z_n \to z$ implies $\sup_{p \in z_n} u(p) \to \sup_{p \in z} u(p)$ for all $u \in U_N$. By dominated convergence, $v(z_n) \to v(z)$ so v is continuous. Now, consider a generic $z \in Z$. For any $p \in \Delta X$, by Lemma 7, we can find p_n with finite support in X^* such that $p_n \to p$. Let $z_n := \{p_n : p \in z\}$ so $z_n \to z$ and $z_n \in \mathcal{K}(\Delta W_n)$ for some finite $W_n \subset X^*$. Define $\alpha_n := v(z_n) \in [0,1]$. Notice that $z_n \sim \alpha_n \delta_{\bar{x}} + (1-\alpha_n)\delta_{\underline{x}}$. Assume $\alpha_n \to \alpha^*$ without loss of generality. Since v is continuous, $v(z) = \alpha^*$. Note that by C1.4, $\underline{x} \preceq z_n \preceq \bar{x}$ for all z_n implies $\underline{x} \preceq z \preceq \bar{x}$. By the standard argument, it can be shown that $z \sim \alpha^* \delta_{\bar{x}} + (1-\alpha^*)\delta_{\underline{x}}$.[73] Finally, to complete the proof, note that $z \succeq y$ if and only if $v(z)\delta_{\bar{x}} + (1 - v(z))\delta_{\underline{x}} \succeq v(y)\delta_{\bar{x}} + (1 - v(y))\delta_{\underline{x}}$ if and only if $v(z) \geq v(y)$. Thus, v represents $\succeq$ on Z.

D.2. **Extension of GP.** Consider a stochastic choice function ρ on Z^f, the finite menus in Z, i.e. ρ_z is a Borel probability measure over z for every $z \in Z^f$. We model ties as in Lu (2016) and Lu and Saito (2018) and let $Z^\circ \subset Z^f$ denote the set of finite menus that contain no ties. We say a probability measure on U_N is *regular* if $u(p) = u(q)$ occurs with probability zero or one for all $p, q \in \Delta X$.

Consider the following conditions: **C2.1** (Monotonicity): $z \subset y$ implies $\rho_z(p) \geq \rho_y(p)$ for all $z, y \in Z$; **C2.2** (Linearity): $\rho_z(p) = \rho_{\alpha z + (1-\alpha)q}(\alpha p + (1-\alpha)q)$ for all $p \in z \in Z$, $\alpha \in (0,1)$, and $q \in \Delta(X)$; **C2.3** (Extremeness): $\rho_z(\mathrm{ext}(z)) = 1$ for any $z \in Z^f$; **C2.4** (Continuity): $\rho : Z^\circ \to \Delta(\Delta X)$ is continuous; **C2.5** (Best-Worst): $\rho(\underline{x}, \bar{x}) = 0$ and $\rho(\bar{x}, x) = \rho(x, \underline{x}) = 1$ for all $x \in X$; **C2.6** (L-continuity): There exists $N > 0$ such that for all $x_1, x_2 \in X$ and $\alpha \in [0,1]$, $|x_1 - x_2| \leq \frac{\alpha}{N}$ implies $\rho\left(\alpha\delta_{\bar{x}} + (1-\alpha)\delta_{x_1}, \alpha\delta_{\underline{x}} + (1-\alpha)\delta_{x_2}\right) = 1$.

Theorem 6 (GP extension)**.** ρ satisfies Condition C2.1–2.6 if and only if there exists a countably additive regular probability measure μ on U_N such that for any $z \in Z^f$ $\rho_z(p) = \mu\{u \in U_N : u(p) \geq u(q) \text{ for all } q \in z\}$.

Other than dealing with ties, the proof is nearly identical to that of Theorem 5. See Section S.2 for details.

[73] Now, suppose $z \succ \alpha\delta_{\bar{x}} + (1-\alpha)\delta_{\underline{x}}$ so we can find some $\beta > \alpha$ such that $z \succ \beta\delta_{\bar{x}} + (1-\beta)\delta_{\underline{x}}$. Since $\alpha_n \to \alpha < \beta$, this means that for large enough n, $\beta\delta_{\bar{x}} + (1-\beta)\delta_{\underline{x}} \succ \alpha_n\delta_{\bar{x}} + (1-\alpha_n)\delta_{\underline{x}} \sim z_n$, where the indifference follows from the representation. By C1.4, we have $\beta\delta_{\bar{x}} + (1-\beta)\delta_{\underline{x}} \succeq z$ yielding a contradiction. The case $z \prec \alpha\delta_{\bar{x}} + (1-\alpha)\delta_{\underline{x}}$ is symmetric.

Appendix E. Proof of Theorem 4 (Representation)

Sufficiency of Axioms: We first prove sufficiency. Note that Axiom 1.1-1.6 corresponds exactly to Conditions C2.1–C2.6 so by Theorem 6, there exists a regular countably additive probability measure μ on U_N s.t. for any $z \in Z^f$, $\rho_z(p) = \mu\{u \in U_N : u(p) \geq u(q)\, \forall q \in z\}$.

Choose any $z_1, z_2 \in Z$. Let $z = \{p_1, p_2\}$ and $y = \{q_1, q_2\}$ where $p_i = \delta_{(c,z_i)}$ and $q_i = \frac{1}{2}p_i \oplus \frac{1}{2}\delta_{(d,w)}$ for $i \in \{1,2\}$. Applying Axiom 2 for $\alpha = \frac{1}{2}$, we have $\mu\{u(p_1) \geq u(p_2)\} = \rho_z(p_1) = \rho_{\frac{1}{2}z+\frac{1}{2}y}\left(\frac{1}{2}p_1 + \frac{1}{2}q_1\right) = \mu\{u(p_1) \geq u(p_2) \text{ and } u(q_1) \geq u(q_2)\}$. Similarly, by Axiom 2, $\mu\{u(q_1) \geq u(q_2)\} = \rho_y(q_1) = \rho_{\frac{1}{2}z+\frac{1}{2}y}\left(\frac{1}{2}p_1 + \frac{1}{2}q_1\right) = \mu\{u(p_1) \geq u(p_2) \text{ and } u(q_1) \geq u(q_2)\}$. Thus, $0 = \mu\{u(p_1) \geq u(p_2) \text{ and } u(q_1) < u(q_2)\} = \mu\{u(p_1) < u(p_2) \text{ and } u(q_1) \geq u(q_2)\}$, so $u(p_1) \geq u(p_2)$ if and only if $u(q_1) \geq u(q_2)$ μ-a.s. For all $c, d \in M$, $z_1, z_2, w \in Z$, and $\lambda \in [0,1]$, we thus have μ-a.s.

(8)
$$u(c, z_1) \geq u(c, z_2) \Leftrightarrow u(\lambda c + (1-\lambda)d, \lambda z_1 + (1-\lambda)w) \geq u(\lambda c + (1-\lambda)d, \lambda z_2 + (1-\lambda)w).$$

Since Z, M and $[0,1]$ are all separable and any $u \in U_n$ is continuous, by the countable additivity of μ, we have that the above holds μ-a.s. for all $c, d \in M$ and $z_1, z_2, w \in Z$ and $\lambda \in [0,1]$. Moreover, we also have that μ-a.s. that for all $c, c' \in M$ and $z_1, z_2, w \in Z$,

(9)
$$u(c, z_1) \geq u(c, z_2) \Leftrightarrow u\left(\frac{1}{2}c + \frac{1}{2}c', \frac{1}{2}z_1 + \frac{1}{2}w\right) \geq u\left(\frac{1}{2}c + \frac{1}{2}c', \frac{1}{2}z_2 + \frac{1}{2}w\right) \Leftrightarrow u(c', z_1) \geq u(c', z_2).$$

We can now define a preference relation $\succeq_u$ on Z for each $u \in U_N$ such that $z \succeq_u y$ if and only if $u(c,z) \geq u(c,y)$. Note that this is well-defined as it does not depend on $c \in M$ by (9) above.

We now show that $\succeq_u$ satisfies Conditions C1.1–1.6 μ-a.s. Note that C1.1 is trivial and C1.3 follows from (8) above. To see C1.2, note that from Axiom 3, for any z, y, if $z \supset y$, then $1 = \rho(z,y) = \rho\left(\delta_{(c,z)}, \delta_{(c,y)}\right) = \mu\{u(c,z) \geq u(c,y)\}$. Since μ is countably additive, $u \in U_N$ is continuous and Z is separable, Condition C1.2 follows. Note that C1.4 follows from the continuity of $u \in U_N$. Finally, by applying Axiom 3 to Axioms 1.5 and 1.6, we obtain C1.5 and C1.6 respectively by the same argument as before. Applying Theorem 5, this means that $\succeq_u$ is represented by $v_u(z) := \int_{U_N} \sup_{p \in z} \tilde{u}(p)\, d\nu_u$, where ν_u is a countably additive probability measure on U_N. Since for every $c \in M$, $u(c, \cdot)$ and v_u represent the same preference, we can write $u(c,z) = \phi_u(c, v_u(z))$, where $\phi_u : M \times [0,1] \to [0,1]$ is strictly increasing in the second argument. Note that this is well-defined as it does not depend on $c \in M$ by (9). The following result shows that μ is the invariant measure of the transition kernel ν_u.

Lemma 8. For any measurable set $B \subset U_N$, $\mu(B) = \int_{U_N} \nu_u(B)\, d\mu$.

Proof. Define the measure μ^* on U_N such that for every measurable $B \subset U_N$, $\mu^*(B) = \int_{U_N} \nu_u(B)\, d\mu$. We will show that $\mu^* = \mu$. Consider finite $z \in Z$ and note that $\rho(z, p_\alpha) = \mu\left\{\sup_{p\in z} u(p) \geq \alpha\right\}$. Thus,

$$\int_{[0,1]} \rho(z, p_\alpha)\, d\alpha = \int_{U_N} \sup_{p\in z} u(p)\, d\mu. \tag{10}$$

Notice also $\rho((c,z),(c,p_\alpha)) = \mu\{\phi_u(c, v_u(z)) \geq \phi_u(c, v_u(p_\alpha))\} = \mu\{v_u(z) \geq \alpha\}$, so

$$\int_{[0,1]} \rho\left(\delta_{(c,z)}, \delta_{(c,p_\alpha)}\right) d\alpha = \int_{U_N} v_u(z)\, d\mu. \tag{11}$$

Applying Axiom 4 to the left-hand sides of (10) and (11), we thus have $\int_{U_N} \sup_{p\in z} u(p)\, d\mu = \int_{U_N} v_u(z)\, d\mu = \int_{U_N}\left(\int_{U_N} \sup_{p\in z} \tilde{u}(p)\, d\nu_u\right) d\mu = \int_{U_N} \sup_{p\in z} u(p)\, d\mu^*$. Letting $z = \{p, p_\alpha\}$, we have $\int_{U_N} \max\{u(p), \alpha\}\, d\mu = \int_{U_N} \max\{u(p), \alpha\}\, d\mu^*$. By Theorem 1.57 of Müller and Stoyan (2002), for any increasing convex function ϕ, $\int_{U_N} \phi(u(p))\, d\mu = \int_{U_N} \phi(u(p))\, d\mu^*$. Since U_N is compact by Lemma 3, $\mu = \mu^*$ by Lemma 4. ■

Let U^1 be the subset of U_N such that for any $u \in U^1$ there exist (ϕ_u, ν_u) satisfying $u(c,z) = \phi_u\left(c, \int_{U_N} \sup_{p\in z} \tilde{u}(p)\, d\nu_u\right)$ so $\mu(U^1) = 1$. Recursively define $U^{n+1} := \{u \in U^n : \nu_u(U^n) = 1\}$ and let $U^* := \bigcap_{n=1}^\infty U^n$. We show that $\mu(U^*) = 1$. First, we show that $\mu(U^n) = 1$ for all n by induction. Suppose $\mu(U^n) = 1$ so by Lemma 8, $1 = \mu(U^n) = \int_{U_N} \nu_u(U^n)\, d\mu$. Thus, $\nu_u(U^n) = 1$ μ-a.s. so $\mu(U^{n+1}) = 1$. Since $\mu(U^1) = 1$, this means that $\mu(U^n) = 1$ for all n. Since $U^{n+1} \subset U^n$, by Proposition 3.6 of Çınlar (2011), $\mu(U^*) = \lim_n \mu(U^n) = 1$.

By Lemma 8 again, we have $1 = \mu(U^*) = \int_{U_N} \nu_u(U^*)\, d\mu$ so $\nu_u(U^*) = 1$ μ-a.s. This means that μ-a.s. that $u(c,z) = \phi_u\left(c, \int_{U^*} \sup_{p\in z} \tilde{u}(p)\, d\nu_u\right)$ and $\rho_z(p) = \mu(B(p,z))$ for any $z \in Z^f$ where $B(p,z) := \{u \in U^* : u(p) \geq u(q)\, \forall q \in z\}$. We can now define a Markov process $[P]$ on $S := U^*$ with invariant distribution μ and transition kernel $P_s := \nu_u$ for all $s = u \in U^*$.

We now prove that the Markov process $[P]$ satisfies Doeblin continuity (i.e., there exists some $\delta > 0$ such that μ-a.s. $\nu_u(A) \geq \delta\mu(A)$ for all measurable A). For this purpose, we will show a density result for support functions. For any $z \in Z$, define the support function $\sigma_z : U_N \to \mathbb{R}$ by $\sigma_z(u) := \sup_{p\in z} u(p)$. Define the sets $\Sigma := \{r(\sigma_z - \sigma_y) : r > 0 \text{ and } z, y \in Z\}$ and $\Sigma^f := \left\{r(\sigma_z - \sigma_y) : r > 0 \text{ and } z, y \in Z^f\right\}$, where σ_z is the support function of $z \in Z$.

Lemma 9. Σ^f is dense in $C(U_N)$.

Proof. Note that for any $z \in Z$, we can find $z_k \in Z_f$ such that $z_k \to z$ (see Lemma 0 of Gul and Pesendorfer (2001)). Thus, $\sigma_{z_k} \to \sigma_z$ by Theorem 7.52 of AB. So Σ^f is dense in Σ. To show the lemma, therefore, it suffices to show that Σ is dense in $C(U_N)$. It is easy to see that Σ is a linear subspace of $C(U_N)$.[74] Given this, by using the argument of Lemma 11 of DLR as well as Stone-Weierstrass Theorem, it is straightforward to show that Σ is dense in $C(U_N)$.[75] ■

Consider any $h \in C(U_N)$ such that $h \geq 0$. By Lemma 9, we can find $h_k \in Z_f$ such that $h_k \to h$. Define $g_k = \max\{h_k, 0\}$ and note that $g_k \to h$ as $h \geq 0$. Moreover, for $h_k = r(\sigma_z - \sigma_y)$ where $z, y \in Z^f$, we have $g_k = r\max\{\sigma_z - \sigma_y, 0\} = r(\sigma_{z\cup y} - \sigma_y) \in \Sigma^f$. By Axiom 5, there exists some $\varepsilon > 0$ s.t. μ-a.s. $\int_{U^*} \sigma_{(1-\varepsilon)(z\cup y)+\varepsilon p_{\varepsilon\bar{y}}} d\nu_u \geq \int_{U^*} \sigma_{(1-\varepsilon)y+\varepsilon p_{\varepsilon\overline{(z\cup y)}}} d\nu_u \Leftrightarrow (1-\varepsilon)\int_{U^*} \sigma_{z\cup y} d\nu_u + \varepsilon\bar{y} \geq (1-\varepsilon)\int_{U^*} \sigma_y d\nu_u + \varepsilon\overline{(z\cup y)} \Leftrightarrow \int_{U^*} (\sigma_{z\cup y} - \sigma_y) d\nu_u \geq \frac{\varepsilon}{1-\varepsilon}\left(\overline{(z\cup y)} - \bar{y}\right) = \delta\int_{U^*} (\sigma_{z\cup y} - \sigma_y) d\mu$,where $\delta := \frac{\varepsilon}{1-\varepsilon}$. Thus, μ-a.s.$\int_{U^*} g_k d\nu_u \geq \delta\int_{U^*} g_k d\mu$. Since $g_k \to h$, this implies that μ-a.s. $\int_{U^*} h d\nu_u \geq \delta\int_{U^*} h d\mu$ by the dominated convergence theorem.

Since U_N is compact, $C(U_N)$ is separable by Lemma 3.99 of AB. Thus, by the countably additivity of μ, μ-a.s. for all nonnegative $h \in C(U_N)$, $\int_{U^*} h d\nu_u \geq \delta\int_{U^*} h d\mu$. Now, by the regularity of ν_u and Urysohn's lemma (Theorem 4.15 of Folland (2013)), for any measurable $A \subset U^*$, there are nonnegative $h_k \in C(U_N)$ such that $h_k \to 1_A$ ν_u-a.s. Thus, by the dominated convergence theorem, μ-.a.s. $P_s(A) = \nu_u(A) = \int_{U^*} \lim_k h_k d\nu_u = \lim_k \int_{U^*} h_k d\nu_u \geq \lim_k \delta\int_{U^*} h_k d\mu = \delta\int_{U^*} \lim_k h_k d\mu = \delta\mu(A)$. Since this implies Doeblin's condition, the Markov process $[P]$ is uniformly ergodic (see Theorem 16.2.3 of Meyn and Tweedie (2012)). By the ergodic theorem, μ-a.s. $\lim_{n\to\infty} \frac{1}{n}\sum_{k\geq 0}^{n} 1_{B(p,z)}(s_k) = \mu(B(p,z)) = \rho(p,z)$ for all $z \in Z^f$ as desired. This concludes the sufficiency proof.

Necessity of Axioms: We now show necessity of the axioms. Note that by Lemma 1, we can consider the ergodic utility process $u_t = u_{s_t}$ with stationary distribution μ. For any $z \in Z^f$, define $B(p,z) := \{u \in U_N : u(p) \geq u(q) \text{ for all } q \in z\}$. By the ergodic theorem, we have for every $z \in Z^f$, $\rho(p,z) = \lim_{n\to\infty} \frac{1}{n}\sum_{k\geq 0}^{n} 1_{B(p,z)}(u_k) = \mu(B(p,z))$. Axioms 1.1-1.6 then follows immediately from Theorem 6.

For Axiom 2, let $p \in z \in Z_c^f$, $y = \lambda z \oplus (1-\lambda)\delta_{(c',z')}$ and $q = \lambda p \oplus (1-\lambda)\delta_{(c',z')} \in y$ where $c, c' \in M$, $z' \in Z$ and $\lambda > 0$. Note that for $p = \delta_{(c,w)}$, $u(p) \geq u(p')$ for all $p' = \delta_{(c,w')} \in z$ if

[74] To see this consider the singleton menu $z = \delta_{\underline{x}}$ and note that by definition, $\sigma_z(u) = u(\underline{x}) = 0$ for all $u \in U_N$. Thus, $0 \in \Sigma$. Next, note that if $r(\sigma_z - \sigma_y) \in \Sigma$, then clearly $\lambda r(\sigma_z - \sigma_y) \in \Sigma$ for all $\lambda \in \mathbb{R}$. Finally, suppose $r_1(\sigma_{z_1} - \sigma_{y_1}), r_2(\sigma_{z_2} - \sigma_{y_2}) \in \Sigma$. Since $r_1, r_2 > 0$, define $\lambda := \frac{r_1}{r_1+r_2}$ so we have $r_1(\sigma_{z_1} - \sigma_{y_1}) + r_2(\sigma_{z_2} - \sigma_{y_2}) = (r_1+r_2)[(\lambda\sigma_{z_1} + (1-\lambda)\sigma_{z_2}) - (\lambda\sigma_{y_1} + (1-\lambda)\sigma_{y_2})]$. Since $\lambda\sigma_{z_1} + (1-\lambda)\sigma_{z_2} = \sigma_{\lambda z_1+(1-\lambda)z_2} \in \Sigma$ (see Lemma 7.54 of AB), we have $r_1(\sigma_{z_1} - \sigma_{y_1}) + r_2(\sigma_{z_2} - \sigma_{y_2}) \in \Sigma$. This shows that Σ is a linear subspace of $C(U_N)$.

[75] Note that for $z = \delta_{\overline{x}}$, $\sigma_z(u) = u(\overline{x}) = 1$ for all $u \in U_N$ so Σ includes constants. That Σ is a vector lattice follows from the same arguments as in Lemma 11 of DLR. Next, we show that Σ separates $C(U_N)$. Choose any $u, v \in U_N$ such that $u \neq v$. Thus, there exists $x \in X$ such that $u(x) \neq v(x)$. If we let $z = \delta_x$, then $\sigma_z(u) = u(x) \neq v(x) = \sigma_z(v)$. Thus, Σ separates $C(U_N)$. Since U_N is compact by Lemma 3, the Stone-Weierstrass Theorem (AB Theorem 9.12) shows that Σ is dense in $C(U_N)$.

and only if $v_u(w) \geq v_u(w')$ for all w' where $v_u(w) := \int_{U_N} \sup_{p \in w} \tilde{u}(p)\, d\nu_u$ and ν_u is the transition kernel corresponding to the ergodic utility process. On the other hand, for all $p' \in z$ and all $q' = \lambda p' \oplus (1-\lambda)\,\delta_{(c',z')} \in y$, $u(q) \geq u(q') \Leftrightarrow u(\lambda c + (1-\lambda)\,c', \lambda w + (1-\lambda)\,z') \geq u(\lambda c + (1-\lambda)\,c', \lambda w' + (1-\lambda)\,z') \Leftrightarrow v_u(\lambda w + (1-\lambda)\,z') \geq v_u(\lambda w' + (1-\lambda)\,z') \Leftrightarrow v_u(w) \geq v_u(w')$ for all w' as $\lambda > 0$. Thus, $u(p) \geq u(p')$ for all $p' \in z$ iff $u(q) \geq u(q')$ for all $q' \in y$. Hence $\rho_z(p) = \mu\{\alpha u(p) + (1-\alpha)\,u(q) \geq \alpha u(p') + (1-\alpha)\,u(q')\, \forall p' \in z, \forall q' \in y\} = \mu\{u(\alpha p + (1-\alpha)\,q) \geq u(\alpha p' + (1-\alpha)\,q')\, \forall p' \in z, \forall q' \in y\} = \rho_{\alpha z + (1-\alpha) y}(\alpha p + (1-\alpha)\,q)$, as desired.

For Axiom 3, suppose $\rho(z,y) = 1$. Let $B := \{u \in U_N : \max_{p \in z} u(p) \geq \max_{q \in y} u(q)\}$ so $\mu(B) = 1$. Since μ is the stationary distribution, $1 = \int_{U_N} \nu_u(B)\, d\mu$ so $\nu_u(B) = 1$ μ-a.s. This implies that $v_u(z) \geq v_u(y)$ μ-a.s. so $\rho\left(\delta_{(c,z)}, \delta_{(c,y)}\right) = 1$ as desired.

For Axiom 4, note that by the same arguments as in Lemma 8, $\int_{[0,1]} \rho(z, p_\alpha)\, d\alpha = \int_{U_N} \sup_{p \in z} u(p)\, d\mu$ and $\int_{[0,1]} \rho\left(\delta_{(c,z)}, \delta_{(c,p_\alpha)}\right) d\alpha = \int_{U_N} v_u(z)\, d\mu = \int_{U_N} \left(\int_{U_N} \sup_{p \in z} \tilde{u}(p)\, d\nu_u\right) d\mu$. The result follows from the fact that μ is the stationary distribution.

Finally, for Axiom 5, suppose $y \subset z$ so $\bar{y} = \int_{U_N} \sup_{p \in y} u(p)\, d\mu \leq \int_{U_N} \sup_{p \in z} u(p)\, d\mu = \bar{z}$. From Lemma 1, we know there exists some δ such that $\nu_u(B) \geq \delta \mu(B)$ for all measurable B so $\int_{U_N} \varphi d\nu_u \geq \int_{U_N} \varphi d\mu$ for all positive measurable functions φ. Let $\varepsilon := \frac{\delta}{1+\delta}$ so $\delta = \frac{\varepsilon}{1-\varepsilon}$. We thus have $v_u((1-\varepsilon)\,z + \varepsilon p_{\bar{y}}) - v_u((1-\varepsilon)\,y + \varepsilon p_{\bar{z}}) = (1-\varepsilon)\,(v_u(z) - v_u(y)) + \varepsilon\,(v_u(p_{\bar{y}}) - v_u(p_{\bar{z}})) = (1-\varepsilon) \int_{U_N} \left(\sup_{p \in z} \tilde{u}(p) - \sup_{p \in y} \tilde{u}(p)\right) d\nu_u + \varepsilon\,(\bar{y} - \bar{z}) \geq (1-\varepsilon)\, \frac{\varepsilon}{1-\varepsilon}\,(\bar{z} - \bar{y}) - \varepsilon\,(\bar{z} - \bar{y}) = 0$. Thus, $\rho\left(\delta_{(c,(1-\varepsilon)z + \varepsilon p_{\bar{y}})}, \delta_{(c,(1-\varepsilon)y + \varepsilon p_{\bar{z}})}\right) = 1$ as desired. This concludes the proof.

References

Aczel, J. (1966): Lectures on functional equations and their applications, Academic Press.

Ahn, D. and T. Sarver (2013): "Preference for Flexibility and Random Choice," Econometrica, 81, 341–361.

Aliprantis, C. and K. Border (2006): Infinite dimensional analysis, Springer.

Alvarez, F. and A. Atkeson (2017): "Random risk aversion and liquidity: a model of asset pricing and trade volumes," Working paper.

Barro, R., J. Fernandez-Villaverde, O. Levintal, and A. Mollerus (2017): "Safe Assets," Working paper.

Block, H. and J. Marschak (1960): "Random Orderings and Stochastic Theories of Response," in Contributions to Probability and Statistics, ed. by I. Olkin, Stanford University Press, 97–132.

Bommier, A., A. Kochov, and F. Le Grand (2017): "On monotone recursive preferences," Econometrica, 85, 1433–1466.

CHEW, S. H. (1989): "Axiomatic utility theories with the betweenness property," Annals of Operation Research, 19, 273–298.

ÇINLAR, E. (2011): Probability and stochastics, vol. 261, Springer Science & Business Media.

DEKEL, E. (1986): "An axiomatic characterization of preferences under uncertainty: weakening the independence axiom," Journal of Economic Theory, 40, 304–318.

DEKEL, E., B. L. LIPMAN, AND A. RUSTICHINI (2001): "Representing preferences with a unique subjective state space," Econometrica, 69, 891–934.

DEKEL, E., B. L. LIPMAN, A. RUSTICHINI, AND T. SARVER (2007): "Representing Preferences with a Unique Subjective State Space: A Corrigendum 1," Econometrica, 75, 591–600.

DILLENBERGER, D., V. KRISHNA, AND P. SADOWSKI (2017): "Subjective Information Choice Processes," Working paper.

DILLENBERGER, D., J. S. LLERAS, P. SADOWSKI, AND N. TAKEOKA (2014): "A theory of subjective learning," Journal of Economic Theory, 153, 287–312.

DURAJ, J. (2018): "Dynamic Random Subjective Expected Utility," Working paper.

EPSTEIN, L. G. (1983): "Stationary cardinal utility and optimal growth under uncertainty," Journal of Economic Theory, 31, 133–152.

EPSTEIN, L. G., E. FARHI, AND T. STRZALECKI (2014): "How much would you pay to resolve long-run risk?" American Economic Review, 104, 2680–97.

EPSTEIN, L. G. AND S. E. ZIN (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," Econometrica, 57, 937–969.

FALMAGNE, J. (1978): "A Representation Theorem for Finite Random Scale Systems," Journal of Mathematical Psychology, 18, 52–72.

FISHBURN, P. (1970): Utility theory for decision making, John Wiley & Sons.

FOLLAND, G. B. (2013): Real analysis: modern techniques and their applications, John Wiley & Sons.

FRICK, M., R. IIJIMA, AND T. STRZALECKI (2018): "Dynamic random utility," Working paper.

FUDENBERG, D. AND T. STRZALECKI (2015): "Dynamic logit with choice aversion," Econometrica, 83, 651–691.

GUL, F. AND W. PESENDORFER (2001): "Temptation and self-control," Econometrica, 69, 1403–1435.

——— (2004): "Self-control and the theory of consumption," Econometrica, 72, 119–158.

——— (2006): "Random expected utility," Econometrica, 74, 121–146.

HIGASHI, Y., K. HYOGO, AND N. TAKEOKA (2009): "Subjective random discounting and intertemporal choice," Journal of Economic Theory, 144, 1015–1053.

Hotz, V. J. and R. A. Miller (1993): "Conditional choice probabilities and the estimation of dynamic models," The Review of Economic Studies, 60, 497–529.

Ke, S. (2018): "Rational expectation of mistakes and a measure of error-proneness," Theoretical Economics, 13, 527–552.

Koopmans, T. C. (1960): "Stationary ordinal utility and impatience," Econometrica, 28, 287–309.

Kreps, D. M. and E. L. Porteus (1978): "Temporal resolution of uncertainty and dynamic choice theory," Econometrica, 185–200.

Krishna, R. V. and P. Sadowski (2014): "Dynamic preference for flexibility," Econometrica, 82, 655–703.

——— (2016): "Randomly Evolving Tastes and Delayed Commitment," Working paper.

Lu, J. (2016): "Random Choice and Private Information," Econometrica, 84, 1983–2027.

Lu, J. and K. Saito (2018): "Random intertemporal choice," Journal of Economic Theory.

Luce, D. (1959): Individual choice behavior, New York: Wiley.

Luce, R. and P. Suppes (1965): "Preference, Utility, and Subjective Probability," in Handbook of Mathematical Psychology, ed. by R. Luce, R. Bush, and E. Galanter, Wiley.

Ma, W. (2018): "Random expected utility theory with a continuum of prizes," Annals of Operations Research.

McFadden, D. (1978): "Modeling the choice of residential location," Transportation Research Record.

——— (1981): "Econometric models of probabilistic choice," Structural analysis of discrete data with econometric applications, 198272.

——— (2001): "Economic choices," American Economic Review, 91, 351–378.

McShane, E. J. (1934): "Extension of range of functions," Bulletin of the American Mathematical Society, 40, 837–842.

Meyn, S. P. and R. L. Tweedie (2012): Markov chains and stochastic stability, Springer Science & Business Media.

Müller, A. and D. Stoyan (2002): Comparison Methods for Stochastic Models and Risks, John Wiley and Sons, Inc.

Rust, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," Econometrica, 999–1033.

——— (1994): "Structural estimation of Markov decision processes," Handbook of econometrics, 4, 3081–3143.

Samuelson, P. (1963): "Risk and Uncertainty: A Fallacy of Large Numbers," Scientia.

Thurstone, L. (1927): "A law of comparative judgment," Psychological Review.

Uzawa, H. (1968): "Time Preference, the Consumption Function, and Optimal Asset Holdings," in Capital and Growth: Papers in Honor of Sir John Hicks, ed. by J. Wolfe, Aldine.

Weil, P. (1990): "Nonexpected Utility in Macroeconomic," Quarterly Journal of Economics, 105, 29–42.

*****Online Supplement (Not for publication)*****

Appendix S.1. Proof of Lemma 2

In this section, we formally define z^t and prove Lemma 2. In order to do so, we first formally define the space of menus following Gul and Pesendorfer (2004). First, define $Z_0 := \{0\}$ and $Z_{t+1} := \mathcal{K}(\Delta(M \times Z_t))$. Also, let $X_{t+1} = \Delta(M \times Z_t)$. Recall that $r_{y,t}(z)$ is the menu that follows $z \in Z$ for t periods and then ends with $y \in Z$ for sure. First, we show that this is well-defined.

Lemma 10. For any $y \in Z$, $r_{y,t} : Z \to Z$ is well-defined.

Proof. We will show by induction that $r_{y,t} : Z \to Z$ is continuous. Clearly this is true for $r_{y,0} = y$. Now, suppose that $r_{y,t-1}$ is continuous so $p_{y,t} \in \Delta X$ is well-defined. We show that $p_{y,t}$ is continuous in $p \in \Delta X$. Consider $p^n \to p$ and let $u : X \to \mathbb{R}$ be continuous and bounded. Note that since $r_{y,t-1}$ is continuous,

$$\int_X u(c,z)\, dp^n_{y,t} = \int_X u(c, r_{y,t-1}(z))\, dp^n \to \int_X u(c, r_{y,t-1}(z))\, dp = \int_X u(c,z)\, dp_{y,t}$$

so $p^n_{y,t} \to p_{y,t}$ as desired. Lemma 1(i) from Gul and Pesendorfer (2004) ensures that $r_{y,t}$ is continuous. Thus, by induction, $r_{y,t}$ is well-defined. ∎

We now extend this notation to menus that end in finite periods, i.e. menus in Z_t. In other words, we will inductively construct the menu $r_{y,t}(z)$ that replicates $z \in Z_i$ for $t \leq i$ periods and ends with $y \in Z_j$ for sure for some j. First, for any $y \in Z_j$, let $r_{y,0}(z) = y$ for any $z \in Z_i$. Given $r_{y,t-1}$, for any $p \in \Delta X_i$ and $t \leq i$, let $p_{y,t} \in \Delta X_{t+j}$ denote the lottery induced by $r_{y,t-1}$, that is, for all measurable $A \times B$, $p_{y,t}(A \times B) = p\left(A \times r^{-1}_{y,t-1}(B)\right)$. Thus, $p_{y,t}$ is the lottery that follows p for $t \leq i$ periods and then yields y for sure. Finally, for any $z \in Z_i$, define $r_{y,t}(z) := \{p_{y,t} : p \in z\}$. In other words, $r_{y,t}(z) \in Z_{t+j}$ is the menu that follows $z \in Z_i$ for $t \leq i$ periods and then ends with y for sure. Note that by the same argument as in Lemma 10, $r_{y,t}$ is also well-defined.

In the following, we define z^t and show that z^t is t-period. If we let $y = 0 \in Z_0$, then $r_{0,t}(z) \in Z_t$ is the t-truncated version of $z \in Z_i$ for $t \leq i$. Following Gul and Pesendorfer (2004), we can now define the space of menus as $Z := \left\{z \in \prod_{t \in T} Z_t \middle| z_t = r_{0,t}(z_{t+1})\right\}$, where z_t denote t-th argument of z for any $t \in T$. We endow Z with the product topology. Theorem A1 of Gul and Pesendorfer (2004) shows that Z is homeomorphic to $\mathcal{K}(\Delta(M \times Z))$.

Given z, we now formally define z^t by constructing a menu $\tilde{z} \in Z$ as follows. First, for any $i \leq t$, let $\tilde{z}_i = z_i$. For $i > t$, set $\tilde{z}_i = r_{\tilde{z}_{i-t},t}(z_t)$ iteratively. Thus, $\tilde{z}$ follows z for $i \leq t$ periods and then replicates itself going forward. Thus $\tilde{z} := (z_1, z_2, \dots, z_t, r_{\tilde{z}_1,t}(z_t), r_{\tilde{z}_2,t}(z_t), \dots) = r_{\tilde{z},t}(z)$. We abuse the notation here and the following; the second equation means $\tilde{z} \in Z$ corresponds to $r_{\tilde{z},t}(z) \in \mathcal{K}(\Delta(M \times Z))$ by the homeomorphism between Z and $\mathcal{K}(\Delta(M \times Z))$. Define $z^t = \tilde{z}$.

We now show that z^t is t-period. We show that $r_{y,t}(Z) \subset R_t(y)$ by induction. First, note that for all $z \in Z$, $r_{y,1}(z) = \{p_{y,1} : p \in z\} \in \mathcal{K}(\Delta(M \times \{y\})) = R_1(y)$, so $r_{y,1}(Z) \subset R_1(y)$. Assume the induction step that $r_{y,t-1}(Z) \subset R_{t-1}(y)$. Thus, for any $p \in \Delta(M \times Z)$, $p_{y,t}(M \times R_{t-1}(y)) \geq p_{y,t}(M \times r_{y,t-1}(Z)) = p(M \times Z) = 1$. Thus, we have $r_{y,t}(z) = \{p_{y,t} : p \in z\} \in R_t(y)$, so $r_{y,t}(Z) \subset R_t(y)$. This shows that $z^t = r_{z^t,t}(z) \in R_t(z^t)$, so z^t is t-period, where the equality means the correspondence based on the homeomorphism. Finally, since $z_i^t = z_i$ for all $i \leq t$, $z^t \to z$ as $i \to \infty$ in the product topology. This concludes the proof of statement (i).

To show statement (ii), fix some finite menu $z \in Z^f$ so from Lemma 2 above, we can find repeated menus z^t such that $z^t \to z$. Since $z^t = r_{z^t,t}(z)$ and z is finite, z^t is also finite by definition. Thus, $z^t \in Z^*$ as desired.

Appendix S.2. Proof of Theorem 6 (GP Extension)

The setup is the same as in Section D. The necessity of the axioms is straightforward. Condition C2.1–2.3 follow from the same arguments as in GP while C2.4 follows from the same argument as in Lu (2016). It is easy to see C2.5 from the representation while C2.6 follows from Lemma 5 above.

We now show the sufficiency and suppose ρ satisfies C2.1–C2.6. The methodology by which we will extend GP parallels the way in which we extended DLR. In particular, the arguments in Lemma 11 below corresponds exactly to those of Lemma 6 in the previous section. The remaining arguments in Lemma 7–13 are new as in GP extension, we need deal with ties.

Lemma 11. There exists a probability measure μ on the Borel σ-algebra corresponding to uniform convergence on U_N such that for all finite $W \subset X^*$ and finite $z \subset \Delta W$, $\rho_z(p) = \mu\{u \in U_N : u(p) \geq u(q) \text{ for all } q \in z\}$.

Proof. We prove this in a series of steps.

Step 1: There exists a probability measure π on the Borel σ-algebra corresponding to pointwise convergence on $\mathbb{R}^{X^*}$ such that for all finite $W \subset X^*$ and finite $z \subset \Delta W$, $\rho_z(p) = \pi\left\{u \in \mathbb{R}^{X^*} : u(p) \geq u(q) \text{ for all } q \in z\right\}$.

Proof. From Gul and Pesendorfer (2006) and Lu (2016), Condition C2.1–C2.4 imply that for each finite $W \subset X^*$ where $\underline{x}, \bar{x} \in W$, there exists a probability measure π_W on $\mathbb{R}^W$ such that for any finite $z \subset \Delta W$, $\rho_z(p) = \pi_W\left\{u \in \mathbb{R}^W : u(p) \geq u(q) \text{ for all } q \in z\right\}$. Moreover, C2.5 implies that we can assume μ-a.s. $0 = u(\underline{x}) \leq u(x) \leq u(\bar{x}) = 1$ for all $x \in X^*$ without loss of generality. By the uniqueness result of GP, all these π_W are consistent.[76] Thus, by Kolmogorov's extension, there exists a measure π on $\mathbb{R}^{X^*}$ such that for all finite $W \subset X^*$ and finite $z \subset \Delta W$, $\rho_z(p) = \pi\left\{u \in \mathbb{R}^{X^*} : u(p) \geq u(q) \text{ for all } q \in z\right\}$. Moreover, we can

[76] Note that this requires normalized utilities.

assume that π is a measure on the Borel σ-algebra corresponding to pointwise convergence on $\mathbb{R}^{X^*}$ (i.e., the product topology, see exercise I.6.35 of Çınlar (2011)). $\square$

By the same argument as in Step 2 of Lemma 6, we have the following:
Step 2: There exists $N > 0$ such that π-a.s. for all $\alpha \in [0,1]$ and $x_1, x_2 \in X^*$, $|x_1 - x_2| \leq \frac{\alpha}{N} \implies \alpha + (1-\alpha) u(x_1) \geq (1-\alpha) u(x_2)$.

By Step 2, Lemma 5 yields $\pi(L_N(X^*)) = 1$. By the Lipschitz version of the Tietze extension theorem (see McShane (1934)), we can extend π on $L_N(X^*)$ to a probability measure μ on $L_N(X)$. By the same argument after Step 2 of Lemma 6, we have $\mu(U_N) = 1$.

Finally, since X is compact, pointwise convergence is equivalent to uniform convergence on U_N. Thus, μ is a measure on the Borel σ-algebra corresponding to uniform convergence. ■

Define $B(p,z) := \{u \in U_N : u(p) \geq u(q) \text{ for all } q \in z\}$, so $B(p,z)$ is μ-measurable. Also define $B(p,q) := B(p,\{p,q\})$ to simplify notation.

We will show that $\rho_z(p) = \mu(B(p,z))$. First, we prove two lemmas which deals with ties in the stochastic choice.

Lemma 12. Suppose $z \in Z^\circ$ and $p_n \to p$ for every $p \in z$ where each p_n has finite support in X^*. If $z_n := \{p_n : p \in z\} \in Z^\circ$, then $\rho_z(p) \leq \mu(B(p,z))$.

Proof. First, note that since $p_n \to p$ for every $p \in z$, $z_n \to z$. Since $z_n, z \in Z^\circ$, Continuity (C2.4) implies that $\rho_z(p) = \lim_n \rho_{z_n}(p_n) = \lim_n \mu(B(p_n, z_n))$, where the last equality follows from the representation as each p_n has finite support in X^*. Note that $\limsup_n 1_{B(p_n,z_n)} \leq 1_{B(p,z)}$. To see why, note that if $\limsup_n 1_{B(p_n,z_n)}(u) = 1$, then there exists a subsequence $\{(p_k, z_k)\}$ such that $u(p_k) \geq u(q_k)$ for all $q_k \in z_k$ so $u(p) \geq u(q)$. Thus, we have

$$\rho_z(p) = \lim_n \int_{U_N} 1_{B(p_n,z_n)} d\mu \leq \int_{U_N} \limsup_n 1_{B(p_n,z_n)} d\mu \leq \int_{U_N} 1_{B(p,z)} d\mu = \mu(B(p,z)),$$

where the first inequality follows from Fatou's Lemma. ■

Lemma 13. (i) If p and q are tied, then $u(p) = u(q)$ a.s.; (ii) If p and q are not tied, then $u(p) \neq u(q)$ a.s.

Proof. First, we show that if p is not tied with $\underline{x}$, then $\rho(\underline{x}, p) = 0$. By Lemma 7, there exists $p_n \to p$ where p_n has finite support in X^*. Let $\tilde{p}_n := \left(1 - \frac{1}{n}\right) p_n + \frac{1}{n}\delta_{\bar{x}}$ and note that $\tilde{p}_n$ cannot be tied with $\underline{x}$ since a.s. $u(\tilde{p}_n) = \left(1 - \frac{1}{n}\right) u(p_n) + \frac{1}{n} > 0$. Note that $\tilde{p}_n \to p$ and each $\tilde{p}_n$ also has finite support in X^*. Since $\{\underline{x}, \tilde{p}_n\} \in Z^\circ$ and $\{\underline{x}, \tilde{p}_n\} \to \{\underline{x}, p\} \in Z^\circ$, Continuity (C2.4) yields $\rho(\underline{x}, p) = \lim_n \rho(\underline{x}, \tilde{p}_n) = \lim_n \mu\{0 \geq u(\tilde{p}_n)\} = 0$, as desired. We now prove the lemma via two steps.
Step 1: If p and q are tied, then $u(p) = u(q)$ a.s.
Proof. First, suppose p is not tied with $\underline{x}$ so $\rho(\underline{x}, p) = 0$ from above. Let $p^\varepsilon := (1-\varepsilon) p + \varepsilon \delta_{\underline{x}}$ so $\rho(p^\varepsilon, p) = 0$ by Linearity (C2.2). Since p and q are tied, $\rho(p^\varepsilon, q) = 0$ by Lemma A.2 of Lu (2016). Consider $z_n^\varepsilon = \{p_n^\varepsilon, q_n\}$ where $p_n^\varepsilon \to p^\varepsilon$, $q_n \to q$ and p_n^ε and q_n both have finite

support in X^* as from Lemma 7. If p_n^ε is tied with q_n, let $\tilde{p}_n^\varepsilon := \left(1-\frac{1}{n}\right)p_n^\varepsilon + \frac{1}{n}\delta_{\underline{x}}$ and $\tilde{q}_n := \left(1-\frac{1}{n}\right)q_n + \frac{1}{n}\delta_{\bar{x}}$, so $\{\tilde{p}_n^\varepsilon, \tilde{q}_n\} \in Z^\circ$. Since $\{\tilde{p}_n^\varepsilon, \tilde{q}_n\} \to \{p^\varepsilon, q\} \in Z^\circ$, by Lemma 12, $1 = \rho(q, p^\varepsilon) \leq \mu(B(q, p^\varepsilon)) = \mu\{u(q) \geq (1-\varepsilon)u(p)\}$. Thus, a.s. $u(p) - u(q) \geq -\varepsilon u(p) \geq -\varepsilon$ for all $\varepsilon > 0$ so $u(q) \geq u(p)$ a.s. By the symmetric reasoning, we have $u(p) \geq u(q)$ a.s. Hence $u(p) = u(q)$ a.s.

Finally, note that if p is tied with $\delta_{\underline{x}}$, then $\frac{1}{2}p + \frac{1}{2}\delta_{\bar{x}}$ is tied with $\frac{1}{2}\delta_{\underline{x}} + \frac{1}{2}\delta_{\bar{x}}$ where the latter is not tied with $\delta_{\underline{x}}$. Applying the above argument yields $\frac{1}{2}u(p) + \frac{1}{2} = \frac{1}{2}$ a.s. or $u(p) = 0$ a.s. as desired. □

Step 2: If p and q are not tied, then $u(p) \neq u(q)$ a.s.

Proof. Let p and q be not tied. Consider $p^\varepsilon := (1-\varepsilon)p + \varepsilon\delta_{\underline{x}}$ and $q^\varepsilon := (1-\varepsilon)q + \varepsilon\delta_{\overline{x}}$ for $\varepsilon > 0$. Note that if p^ε and q^ε are tied, then from (i), we have a.s. $u(p) = u(q) + \frac{\varepsilon}{1-\varepsilon}$. Thus, we can choose $\varepsilon \to 0$ such that p^ε and q^ε are not tied. Consider $z_n^\varepsilon = \{p_n^\varepsilon, q_n^\varepsilon\}$ where $p_n^\varepsilon \to p^\varepsilon$, $q_n^\varepsilon \to q^\varepsilon$ and p_n^ε and q_n both have finite support in X^* as above. Again, let $\tilde{p}_n^\varepsilon := \left(1-\frac{1}{n}\right)p_n^\varepsilon + \frac{1}{n}\delta_{\underline{x}}$ and $\tilde{q}_n^\varepsilon := \left(1-\frac{1}{n}\right)q_n^\varepsilon + \frac{1}{n}\delta_{\bar{x}}$, so $\{\tilde{p}_n^\varepsilon, \tilde{q}_n^\varepsilon\} \in Z^\circ$. Since $\{\tilde{p}_n^\varepsilon, \tilde{q}_n^\varepsilon\} \to \{p^\varepsilon, q^\varepsilon\} \in Z^\circ$, by Lemma 12, $\rho(p^\varepsilon, q^\varepsilon) \leq \mu(B(p^\varepsilon, q^\varepsilon)) = \mu\left\{u(p) - u(q) \geq \frac{\varepsilon}{1-\varepsilon}\right\}$. As $\varepsilon \searrow 0$, $\{p^\varepsilon, q^\varepsilon\} \to \{p, q\} \in Z^\circ$ so by Continuity (C2.4), $\rho(p,q) = \lim_{\varepsilon\searrow 0}\rho(p^\varepsilon, q^\varepsilon) \leq \lim_{\varepsilon\searrow 0}\mu\left\{u(p) - u(q) \geq \frac{\varepsilon}{1-\varepsilon}\right\} = \mu\{u(p) > u(q)\}$. By symmetric reasoning, we have $\rho(q,p) \leq \mu\{u(p) > u(q)\}$, so $1 = \rho(p,q) + \rho(q,p) \leq \mu\{u(p) > u(q)\} + \mu\{u(p) > u(q)\}$. Thus, $u(p) = u(q)$ has μ-measure zero. □

■

We now complete the proof of Theorem 6. Let $z \in Z^\circ$ and $p_n \to p$ for every $p \in z$ where each p_n has finite support in X^*. Note that $z_n := \{p_n : p \in z\} \to z$. Suppose there exists an infinite subsequence such that $z_n \notin Z^\circ$. Thus, there must be a subsequence $p_n, q_n \in z_n$ that are tied for each n. By Lemma 13, $u(q_n) = u(p_n)$ a.s. so $u(q) = u(p)$ a.s. By Lemma 13 again, this means p and q are tied, contradicting $z \in Z^\circ$. Thus, we can assume that $z_n \in Z^\circ$ so by Lemma 12, we have $\rho_z(p) \leq \mu(B(p,z))$.

Finally, let $z_0 \subset z$ be such that $z_0 \in Z^\circ$ so $\rho_{z_0}(p) \leq \mu(B(p, z_0))$. Suppose $\rho_{z_0}(p) < \mu(B(p, z_0))$ for some $p \in z_0$. Thus, $1 = \sum_{p\in z_0}\rho_{z_0}(p) < \sum_{p\in z_0}\mu(B(p,z_0)) \leq 1$, where the last inequality follows from Lemma 13 and the fact that z_0 has no ties. Since this yields a contradiction, it must be that $\rho_{z_0}(p) = \mu(B(p, z_0))$ for all $p \in z_0$. Now, for any $p \in z$, we can find some $p_0 \in z_0$ tied with p. By Lemma A.2 from Lu (2016), we have $\rho_z(p) = \rho_{z_0}(p_0) = \mu(B(p_0, z)) = \mu(B(p, z))$, as desired.

Appendix S.3. Proof of Lemma 5

Suppose (i) is true. Fix some $\bar{\alpha} < 1$ and consider $x_1, x_2 \in X^*$. First suppose $|x_1 - x_2| N = \alpha \leq \bar{\alpha} < 1$. We thus have $\alpha v(\underline{x}) + (1-\alpha)v(x_2) \leq \alpha v(\bar{x}) + (1-\alpha)v(x_1)$.

Hence, $v(x_2) - v(x_1) \le \frac{\alpha}{1-\alpha} = \frac{N}{1-\alpha}|x_1 - x_2| \le \frac{N}{1-\bar{\alpha}}|x_1 - x_2|$. Now suppose $|x_1 - x_2| N = \alpha > \bar{\alpha}$.

Since X is a convex metric space, we can find $y_i := \left(1 - \frac{i}{n}\right) x_1 + \frac{i}{n} x_2 \in X$ for $i \in \{0, 1, \dots, n\}$ such that $|y_{i+1} - y_i| = \frac{1}{n}|x_1 - x_2| < \frac{\bar{\alpha}}{N}$. Since X^* is dense in X and the metric mapping is continuous, we can choose n large enough such that for each $\varepsilon > 0$, we can find $y_i^* \in X^*$ such that $|y_i - y_i^*| \le \varepsilon$ and $\left|y_{i+1}^* - y_i^*\right| < \frac{\bar{\alpha}}{N}$ for all i. From the argument above, we have $v\left(y_{i+1}^*\right) - v\left(y_i^*\right) \le \frac{N}{1-\bar{\alpha}}\left|y_{i+1}^* - y_i^*\right| \le \frac{N}{1-\bar{\alpha}}\left(|y_{i+1} - y_i| + \left|y_{i+1}^* - y_{i+1}\right| + |y_i^* - y_i|\right) \le \frac{N}{1-\bar{\alpha}}\left(|y_{i+1} - y_i| + 2\varepsilon\right) = \frac{N}{1-\bar{\alpha}}\left(\frac{1}{n}|x_1 - x_2| + 2\varepsilon\right)$.

Since we can let $y_0^* = y_0 = x_1$ and $y_n^* = y_n = x_2$, this implies that $v(x_2) - v(x_1) \le \sum_{1 \le i \le n}\left|v\left(y_i^*\right) - v\left(y_{i-1}^*\right)\right| \le \frac{N}{1-\bar{\alpha}}\left(|x_1 - x_2| + 2n\varepsilon\right)$. Taking $\varepsilon \to 0$ yields $v(x_2) - v(x_1) \le \frac{N}{1-\bar{\alpha}}|x_1 - x_2|$. Since $\frac{N}{1-\bar{\alpha}} \to N$ as $\bar{\alpha} \to 0$, this means that $|v(x_2) - v(x_1)| \le N|x_1 - x_2|$ for all $x_1, x_2 \in X^*$. Thus, v is Lipschitz continuous with bound N as desired.

Now, suppose (ii) is satisfied. Note that if $\alpha = 1$, then the result is trivial so assume $\alpha < 1$. Suppose that $|x_1 - x_2| \le \frac{\alpha}{N}$ and since $v \in L_N(X^*)$, $v(x_2) - v(x_1) \le N|x_1 - x_2| \le \frac{N}{1-\alpha}|x_1 - x_2| \le \frac{\alpha}{1-\alpha}$. Rearranging yields $\alpha v(\underline{x}) + (1-\alpha) v(x_2) \le \alpha v(\bar{x}) + (1-\alpha) v(x_1)$, as desired.

Appendix S.4. Stochastic Epstein-Zin and RI

Under stochastic Epstein-Zin, non-standard intertemporal preferences manifest themselves in spurious violations of the classic independence axiom. Recall from Theorem 3 that RI along with IRU characterize ICM. For an Epstein-Zin agent, PEU (i.e., $\psi_s \le RRA_s$) or PLU (i.e., $\psi_s \ge RRA_s$) can be detected by how RI is violated. Let $\ge_{FOSD}$ denote first-order stochastic dominance.

Proposition 5. Suppose ρ is stochastic Epstein-Zin. For 1-period $z \in Z^*$ and $p_1 \ge_{FOSD} r$ for all $p \in z$, (i) $\psi_s \le RRA_s$ a.s. implies $\rho_z\left(\delta_{(c,z)}\right) \le \rho_{az \otimes (1-a) r}\left(a\delta_{(c,z)} \otimes (1-a) r\right)$; (ii) $\psi_s \ge RRA_s$ a.s. implies $\rho_z\left(\delta_{(c,z)}\right) \ge \rho_{az \otimes (1-a) r}\left(a\delta_{(c,z)} \otimes (1-a) r\right)$.

Proof. First, suppose $\psi_s \le RRA_s$ a.s. Let $y = az \otimes (1-a) r$. Since $p_1 \ge_{FOSD} r$ for all $p \in z$, $v_{s_t}(z) = \mathbb{E}_{s_t}\left[\sup_{q \in z} u_{s_{t+1}}(q)\right] \ge \mathbb{E}_{s_t}\left[\sup_{q \in y} u_{s_{t+1}}(q)\right] = v_{s_t}(y)$. Let $v_2 = v_s(z)$ and $v_1 = v_s(z)$ so $v_2 \le v_1$. Now, for any $p \in z$, $u_s\left(\delta_{(c,z)}\right) \ge u_s(p) \Leftrightarrow \phi_s(c, v_1) \ge \int_M \phi_s(d, v_1)\, dp_1$. On the other hand, $u_s\left(a\delta_{(c,z)} \otimes (1-a) r\right) \ge u_s(ap \otimes (1-a) r) \Leftrightarrow a u_s(c, y) + (1-a)\int_M u_s(c', y)\, dr$ $\ge a\int_M u_s(c', y)\, dp_1 + (1-a)\int_M u_s(c', y)\, dr \Leftrightarrow \phi_s(c, v_2) \ge \int_M \phi_s(c', v_2)\, dp_1$

Since $\psi_s \le RRA_s$, $\phi_s(\cdot, v_1)$ is more convex than $\phi_s(\cdot, v_2)$ as in the proof of Proposition 1. Thus, for every $p \in z$, $u_s\left(\delta_{(c,z)}\right) \ge u_s(p)$ implies $u_s\left(a\delta_{(c,z)} \otimes (1-a) r\right) \ge u_s(ap \otimes (1-a) r)$ so the conclusion follows. The case for $\psi_s \ge RRA_s$ a.s. is symmetric. ■

Proposition 5 illustrates the type of permissible violation of the classic independence axiom in the repeated choice setup. For example, under strict PEU, if z consists of a risky and a safe option, then the probability of choosing the safe option will strictly increase if we mix all options with the worst consumption. Note that the act of mixing changes the agent's continuation value; when intertemporal preferences are non-standard as in Epstein-Zin, this generates violations of repeated independence. We can interpret this as a spurious violation of the independence axiom due to ignoring the intertemporal structure of the problem.

Note that this does not permit *any* violation of independence; for example, the agent will never strictly prefer mixtures. This is because the agent is still an expected utility maximizer on the larger outcome space of pairs of consumption and continuation menus.

Let $p_1, q_1 \in \Delta(M)$. Given a repeated menu $z = \{(p_1, z), (q_1, z)\}$, the agent will never choose the mixture $\left(\frac{1}{2}p_1 + \frac{1}{2}q_1, y\right)$ in the repeated menu $y = \left\{(p_1, y), \left(\frac{1}{2}p_1 + \frac{1}{2}q_1, y\right), (q_1, y)\right\}$. Even though there may be consumption smoothing due to intertemporal preferences, a stochastic Epstein-Zin agent will never exhibit a strict preference for ex-ante hedging; in other words, our model satisfies the stochastic version of betweenness from Dekel (1986) and Chew (1989).

Let $r = \delta_0$ and note that for any 1-period $z \in Z$, $a\delta_{(c,z)} \otimes (1-a)\,\delta_0 \to \delta_{\underline{x}}$ as $a \to 0$. This suggests the following comparative statics result.

Proposition 6. Suppose ρ and ρ' are both stochastic Epstein-Zin with respective risk aversion distributions π_{RRA} and π'_{RRA}. Then $\pi_{RRA} \geq_{FOSD} \pi'_{RRA}$ iff for all 1-period $z \in Z^*$, $\lim_{a\to 0} \rho_{az\otimes(1-a)\delta_0}\left(a\delta_{(c,z)} \otimes (1-a)\,\delta_0\right) \leq \lim_{a\to 0} \rho'_{az\otimes(1-a)\delta_0}\left(a\delta_{(c,z)} \otimes (1-a)\,\delta_0\right)$.

Proof. Let $z = \left\{\delta_{(c,z)}, p\right\}$ and $y_a = az \otimes (1-a)\,\delta_0$. Note that $y_a \to \delta_{\underline{x}}$ as $a \to 0$ so $\lim_{a\to 0} v_{s_t}(y_a) = v_{s_t}(\underline{x}) = 0$. Let $w_s(c) = c^{1-RRA_s}$ denote CRRA utility. We thus have $\lim_{a\to 0} \rho_{az\otimes(1-a)r}\left(a\delta_{(c,z)} \otimes (1-a)\,\delta_0\right) = \lim_{a\to 0} \pi\left\{\phi_s(c, v(y_a)) \geq \int_M \phi_s(c', v(y_a))\,dp_1\right\}$ $= \pi\left\{\phi_s(c, 0) \geq \int_M \phi_s(c', 0)\,dp_1\right\} = \pi_{RRA}\left\{w_s(c) \geq w_s(p_1)\right\}$. The conclusion follows from the fact that $\pi_{RRA} \geq_{FOSD} \pi'_{RRA}$ iff $\pi_{RRA}\left\{w_s(c) \geq w_s(p_1)\right\} \leq \pi'_{RRA}\left\{w_s(c) \geq w_s(p_1)\right\}$ for all $c \in M$ and $p_1 \in \Delta M$. ■