

Dynamic Random Utility*

Mira Frick

Ryota Iijima

Tomasz Strzalecki

Abstract

We provide an axiomatic analysis of dynamic random utility, characterizing the stochastic choice behavior of agents who solve dynamic decision problems by maximizing some stochastic process (U_t) of utilities. We show first that even when (U_t) is arbitrary, dynamic random utility imposes new testable restrictions on how behavior *across periods* is related, over and above period-by-period analogs of the static random utility axioms: An important feature of dynamic random utility is that behavior may appear *history dependent*, because past choices reveal information about agents' past utilities and (U_t) may be serially correlated; however, our key new axioms highlight that the model entails specific *limits* on the form of history dependence that can arise. Second, we show that when agents' choices today influence their menu tomorrow (e.g., in consumption savings or stopping problems), imposing natural Bayesian rationality axioms restricts the form of randomness that (U_t) can display. By contrast, a specification of utility shocks that is widely used in empirical work violates these restrictions, leading to behavior that may display a negative option value and can produce biased parameter estimates. Finally, dynamic stochastic choice data allows us to characterize important special cases of random utility—in particular, learning and taste persistence—that on static domains are indistinguishable from the general model.

*This version: 2 November 2018. Frick: Yale University (mira.frick@yale.edu); Iijima: Yale University (ryota.ijima@yale.edu); Strzalecki: Harvard University (tomasz_strzalecki@harvard.edu). This research was supported by National Science Foundation grant SES-1255062. We thank a co-editor and four anonymous referees for detailed and insightful comments that greatly improved the paper. We also thank David Ahn, Jose Apesteguia, Miguel Ballester, Dirk Bergemann, Francesco Cerigioni, Jetlir Duraj, Drew Fudenberg, Myrto Kalouptsi, Emir Kamenica, Shaowei Ke, Daria Khromenkova, Yves Le Yaouanq, Jay Lu, Yusuke Narita, Andriy Norets, Ariel Pakes, Larry Samuelson, Jesse Shapiro, Michael Whinston, as well as numerous seminar audiences. A joint file, including both the main text and supplementary appendix, is available at https://drive.google.com/open?id=1JIrSyzkpi1-0yNfDYoQ_dc3Se6qLBeDM

1 Introduction

1.1 Motivation

Random utility models are widely used throughout economics. In the static model, the agent chooses from her choice set by maximizing a random utility function U . In the dynamic model, the agent solves a dynamic decision problem, subject to a stochastic process (U_t) of utilities. The key feature of the model is an informational asymmetry between the agent (who knows her realized utility) and the analyst (who does not). In both the static and dynamic setting, this asymmetry gives rise to choice behavior that appears stochastic to the analyst but is deterministic from the point of view of the agent.¹

A classic literature in decision theory axiomatically characterizes the stochastic choice behavior that is implied by any *static* random utility model, regardless of the details of the agent’s random utility function (see Section 7.1). Axiomatic analysis helps shed light on which behavior (e.g., the “attraction effect”) this model rules out, as well as which behavior (e.g., the restrictive “independence of irrelevant alternatives” assumption) is implied only by specific parametric versions of random utility but not by the general model.² Moreover, some axioms have inspired empirical tests of the model (e.g., Hausman and McFadden, 1984; Kitamura and Stoye, 2018).

This paper provides the first axiomatic characterization of the fully general and non-parametric model of *dynamic* random utility. Our analysis yields the following main insights. First, we show that even when the agent’s utility process is arbitrary, dynamic random utility imposes new testable restrictions on how behavior *across periods* is related, over and above period-by-period analogs of the static random utility axioms. An important feature of dynamic random utility is that behavior generally appears *history dependent* to the analyst, because past choices reveal some information about past utilities and (U_t) may display serial correlation: For example, we expect an agent’s probability of voting Republican in 2020 to be different conditional on voting Republican in 2016 than conditional on voting Democrat in 2016, as her past voting behavior reflects her past political preferences, which are typically at least somewhat persistent.³ However, our key new axioms highlight that any dynamic random utility model imposes specific *limits* on the form of history dependence that can arise.

Second, in many dynamic decision problems, such as consumption-savings or optimal stop-

¹We interpret this stochastic choice data as the analyst’s observation of a large *population* of individuals whose heterogeneous (resp. stochastically evolving) utilities are realized according to U (resp. U_t). By convention, we use “the agent” to refer to any one of these individuals whose identity is unknown to the analyst; see Section 2.2.3.

²See, e.g., Huber, Payne, and Puto (1982) and Block and Marschak (1960).

³Throughout the paper, we restrict attention to the case where utilities U_t evolve *exogenously*. Thus, from the point of view of the agent, past choices have no effect on current utility. As we discuss in Section 7.2, our characterization can be extended to allow for the latter effect (e.g., due to habit formation or active learning).

ping problems, the agent’s choices today also influence her menu tomorrow. Our second main result shows that imposing natural Bayesian rationality axioms on behavior in such settings restricts the random evolution of the agent’s utility process. Specifically, randomness in U_t must arise from shocks to the agent’s evaluation of instantaneous *consumptions*, and utilities across periods are related by a Bellman equation that correctly anticipates future shocks. By contrast, we show that a second form of utility shocks, *shocks to actions*, that are statistically convenient and widely used in the empirical literature on dynamic discrete choice (DDC) can give rise to behavior that violates basic features of Bayesian rationality.

Third, even when the agent only faces sequences of static choice problems and today’s choices do not influence tomorrow’s menu, dynamic stochastic choice data makes it possible to distinguish important models of utility shocks that are indistinguishable on static domains. In particular, relative to the case of arbitrarily evolving utilities, we characterize the additional behavioral content of an agent with a fixed but unknown utility about which she *learns* over time and of an agent who displays *taste persistence*.

Our results are complementary to the DDC literature. The latter studies dynamic random utility models, and associated phenomena such as history dependence and choice persistence, with focus on identification and estimation. This paper provides decision-theoretic foundations that focus on testable implications, comparative statics, and distinctions between key special cases of dynamic random utility. We hope that the modeling tradeoff between shocks to consumption and shocks to actions that we highlight will stimulate a conversation about desirable properties of models and the ways to resolve this tradeoff.

1.2 Overview

Section 2 sets up our model of *dynamic random expected utility (DREU)*. This generalizes the static random expected utility framework of Gul and Pesendorfer (2006) to decision trees as defined by Kreps and Porteus (1978). Each period t , the agent chooses from a menu A_t of lotteries p_t that determine both her current consumption z_t and tomorrow’s menu A_{t+1} . Her choices maximize a random vNM utility U_t whose realizations are governed by a probability distribution μ over a state space Ω that allows for arbitrary serial correlation of utilities. From the point of view of the analyst, this generates a history dependent stochastic choice rule: A *history* $h^{t-1} = (A_0, p_0, \dots, A_{t-1}, p_{t-1})$ summarizes that the agent chose lottery p_0 from menu A_0 , then faced A_1 and chose p_1 , and so on. Following any history h^{t-1} , the analyst observes the *conditional choice probability* $\rho_t(p_t; A_t | h^{t-1})$ of p_t from menu A_t . In particular, in period $t = 0$ there is no history to condition on, so ignoring ties, $\rho_0(p_0; A_0) = \mu(U_0(p_0) = \max_{q_0 \in A_0} U_0(q_0))$,

just as under static random utility. In period $t = 1$, we have

$$\rho_1(p_1; A_1 | A_0, p_0) = \mu \left(U_1(p_1) = \max_{q_1 \in A_1} U_1(q_1) \mid U_0(p_0) = \max_{q_0 \in A_0} U_0(q_0) \right),$$

and analogously for any $t > 1$.

Section 3 characterizes DREU. Our key new axioms capture the following idea: As history dependence under DREU results purely from the information that past choices reveal about the agent’s utility, this entails certain forms of history *independence*. Specifically, we identify two simple cases in which histories h^{t-1} and g^{t-1} reveal the *same* information about the agent’s utilities, and we require that choice behavior $\rho_t(\cdot | h^{t-1})$ and $\rho_t(\cdot | g^{t-1})$ following two such histories must coincide. Axiom 1, *contraction history independence*, considers the case where h^{t-1} can be obtained from g^{t-1} by eliminating some options that are “irrelevant” to choices along the history g^{t-1} (see Example 1 for an illustration). This rules out certain dynamically “irrational” behavior such as the “mere exposure effect,” where the mere presence of some option that the agent does not choose today might affect her behavior tomorrow.

Axiom 2, *linear history independence*, considers h^{t-1} and g^{t-1} that are “linear combinations” of each other. As Example 2 illustrates, this axiom provides a conceptual justification for a lottery-based extrapolation procedure we use to overcome the “limited observability” problem, an important challenge specific to the dynamic setting: Whereas in the static domain the analyst observes choices from all possible menus, in the dynamic setting any history of past choices restricts the set of current and future choice problems, which over time, severely limits the history-dependent choice data observable to the analyst. Theorem 1 shows that Axioms 1 and 2, along with a continuity condition and Gul and Pesendorfer’s (2006) axioms that ensure *static* random utility maximization at each history, fully characterize DREU.

In DREU, the utility process (U_t) is unrestricted and in principle allows the agent to be myopic or suffer from temptation problems. Section 4 studies the important special case of *Bayesian evolving utility (BEU)*, where the agent is dynamically sophisticated and forward-looking with a correct assessment of option value. BEU is obtained by imposing Bayesian rationality axioms on DREU; specifically, we adapt the preference for flexibility and dynamic sophistication conditions from the menu preference literature to our stochastic choice setting. Theorem 2 shows that these axioms yield a utility process (U_t) where the agent’s evaluation of current consumption z_t and continuation menu A_{t+1} satisfies the Bellman equation

$$U_t(z_t, A_{t+1}) = u_t(z_t) + \delta_t \mathbb{E} \left[\max_{p_{t+1} \in A_{t+1}} U_{t+1}(p_{t+1}) \mid \mathcal{F}_t \right]$$

for some process (u_t) of random felicities, (δ_t) of stochastic discount factors, and a filtration (\mathcal{F}_t) that represents the agent’s private information.

Section 5 contrasts BEU with dynamic discrete choice (DDC) models. BEU is a special case of the most general DDC model. However, for estimation purposes, most DDC models subject the agent’s utility over continuation menus to additional randomness (*shocks to actions*) that may be completely detached from their continuation value; Example 3 illustrates this in the context of a simple stopping problem. Relative to BEU, we highlight the following modeling tradeoff. On the one hand, shocks to actions are statistically more convenient, but unlike BEU, they can lead to violations of a key feature of Bayesian rationality, positive option value: For example, we show that more often than not, the agent chooses to make decisions as *early* as possible, even when delay is costless and could provide her with better information about her payoffs; moreover, greater uncertainty about her utilities may lead her to value delay *less*. In settings such as Example 3, we also show that the conceptual differences between the two models translate into systematically different parameter estimates.

Finally, Section 6 restricts to the simpler subdomain of *atemporal consumption problems*, where each period agents choose only (lotteries over) today’s consumption and their current choices do not affect tomorrow’s menu. Choice data on this domain is often featured in empirical work (e.g., the literature on brand choice dynamics) and an important regularity is that choices tend to display some “persistence.” As Example 1 illustrates, we show that two natural forms of choice persistence capture the additional behavioral content of two important special cases of BEU: *Bayesian evolving beliefs (BEB)*, where current felicity u_t represents the agent’s expectation of her fixed but unknown tastes \tilde{u} about which she receives new information each period; and the case where u_t displays a non-parametric form of *taste persistence*. On our original domain, Theorem 3 provides an alternative characterization of BEB in terms of a consumption stationarity axiom that reflects the martingale property of beliefs. We also show that, unlike BEU, under BEB the agent’s discount factor process is uniquely identified.

1.3 Illustrative Examples

Example 1 (Brand choice dynamics). A large marketing literature studies repeated consumer choices between different brands.⁴ In this data, *history-dependent* choices are widely observed; as an illustration, in Figure 1 (left), brand x is most popular at all nodes, but period 1 behavior differs substantially across consumers who chose x in period 0 and those who chose y .

As discussed in the introduction, under dynamic random utility, history dependence can result from the fact that agents’ tastes (u_t) may be serially correlated. However, our axioms in Section 3.1 show that even under arbitrary serial correlation of utilities, there are limits on the forms of history dependence that can arise. For example, suppose an ex ante identical population of consumers additionally face brand z in period 0 and choice frequencies are as in

⁴See the references in Section 6.




Figure 1: Brand choice.

Figure 1 (right). As we will see, our Axiom 1 (contraction history independence) implies that the period-1 choice frequencies among consumers who chose x in period 0 must be the *same* in both decision trees in Figure 1. This is because z is an “irrelevant” alternative from the point of view of x , as it does not affect x ’s demand share.

In addition, we characterize precisely which non-parametric forms of serial correlation in (u_t) correspond to certain widely documented forms of history dependence. Specifically, the data in Figure 1 (left) displays *consumption inertia*, where a sizable share of consumers who chose y in period 0 chooses it again over x in period 1, and *consumption persistence*, where the share of consumers choosing y in period 1 is higher among those who chose y in period 0 than among those who chose x in period 0. Section 6 shows that on simple domains such as the one in Figure 1, consumption inertia characterizes consumers with fixed but unknown utilities \tilde{u} about which they *learn* over time; i.e., u_t represents their expectations of \tilde{u} given period t information (as in our BEB model). By contrast, consumption persistence characterizes consumers whose tastes u_t display a particular form of positive serial correlation that we call *taste persistence*. We also provide comparative statics of behavior with respect to the amount of taste persistence. ▲

Example 2 (School choice). Unlike Example 1, in many economic settings agents’ choices today also affect their menus tomorrow. Figure 2 (left) provides a stylized example in the context of school choice. In period 0, parents decide to enroll their child in one of two elementary schools, which differ along many decision-relevant dimensions. Upon enrolling, parents must then choose between a number of after-school care options: H (stay at home/leave the child with relatives); P (a high quality but high cost private after-school center); or S (a more basic and lower cost after-school program offered *only* by school 1). Thus, choosing school 1 leads to




Figure 2: School choice.

period-1 menu $\{H, P, S\}$, whereas school 2 leads to menu $\{H, P\}$.

In such settings, history-dependent choice behavior can result from dynamic selection effects: Different types of parents select into each school, so the observed choices from $\{H, P, S\}$ and from $\{H, P\}$ do not reflect the unconditional choice frequencies that would arise if all parents made choices from either menu.⁵ Failure to account for this may lead to *spurious violations of random utility*. For example, in Figure 2 (left), the share of parents choosing P is larger at school 1 (30%) than school 2 (20%), despite the fact that more options are available at school 1. Ignoring history dependence, this behavior appears to violate the Regularity axiom, which is a well-known implication of static random utility (Block and Marschak, 1960). However, it is entirely consistent with *dynamic* random utility maximization, because under serially correlated private information the preferences of parents at each school will differ.⁶ In Section 3.3, we show that under dynamic random utility, period-by-period versions of the static random utility axioms are valid only if the analyst controls for past choices.

As discussed in the introduction, another important challenge implied by history dependence is *limited observability*. For example, in the left-hand decision tree in Figure 2 we do not observe the counterfactual frequencies with which parents at school 1 would choose between H and P if S was not available to them; and given dynamic selection, we cannot simply infer these from the corresponding choice frequencies of parents at school 2. However, in practice many schools

⁵This is a key difference between our setting and (i) Ahn and Sarver (2013) and (ii) Fudenberg and Strzalecki (2015): (i) assume that period-0 choices between menus are *deterministic*; (ii) assume that the agent's utility process is *i.i.d.* In either case, there are no dynamic selection effects and period-1 choices from menus are *history-independent*.

⁶E.g., a preference for other features of school 2 may happen to be strongly correlated with a preference for H ; or parents for whom H is more costly might select disproportionately into school 1 because it expands their outside-the-home options.

ration their seats via lotteries, a fact that is widely exploited in the empirical literature on school choice to generate quasi-experimental variation.⁷ This is illustrated in the right-hand tree in the figure, where each application to school 1 is successful with probability λ and the parent must select school 2 otherwise. In Section 3.2, we show how in such settings, the analyst can extrapolate the choices that school 1 parents in the left-hand tree would make from the set $\{H, P\}$ by looking at choices of parents in the right-hand tree who applied to school 1 but were rejected by the lottery. Our Axiom 2 (linear history independence) provides a conceptual justification for this extrapolation procedure, as it implies that the inference does not depend on the randomization probability λ . \blacktriangle

Example 3 (Optimal stopping). Consider the following optimal stopping problem. The agent can consume a in period 0 (and nothing in period 1) or defer consumption and then choose between a or b in period 1, whichever she prefers at that point. For example, suppose b is a more expensive substitute for a that the agent can only afford by foregoing consumption in period 0 and accumulating enough savings by period 1; or b is a new model with release date scheduled for period 1. Figure 3 depicts the decision tree, where $A_1 := \{a, b\}$ and $A_0 := \{a, A_1\}$.




Figure 3: Optimal stopping.

How the agent resolves the tradeoff between immediate consumption and the option value of delay depends on the underlying structural parameters: the distribution of utility shocks and the discount factor δ . Section 5 contrasts two models of utility shocks: *Shocks to consumption* apply only to instantaneous consumptions and affect the agent’s evaluation of tomorrow’s menu only through her anticipation of future shocks to consumption; by contrast, *shocks to actions* subject today’s evaluation of tomorrow’s menu to an additional shock that may be completely detached from its continuation value. We show that this is the main difference between our BEU model (shocks to consumption) and many widely used models in the DDC literature (shocks to actions).

In the present example, compare the following specifications of BEU and DDC, where all shocks are assumed i.i.d. for simplicity. BEU assigns shocks ε_0^a , ε_1^a , and ε_1^b to the instantaneous

⁷E.g., Abdulkadiroglu, Angrist, Narita, and Pathak (forthcoming); Angrist, Hull, Pathak, and Walters (forthcoming); Deming (2011); Deming, Hastings, Kane, and Staiger (2014).

consumptions in periods 0 and 1, and the latter two enter the continuation value to menu A_1 :

$$U_0^{\text{BEU}}(a) = v(a) + \varepsilon_0^a \text{ and } U_0^{\text{BEU}}(A_1) = \delta \mathbb{E} [\max\{v(a) + \varepsilon_1^a, v(b) + \varepsilon_1^b\}].$$

By contrast, i.i.d. DDC assigns an additional shock $\varepsilon_0^{A_1}$ to the period 0 action of delaying and choosing menu A_1 , even though this entails no instantaneous consumption:

$$U_0^{\text{DDC}}(a) = v(a) + \varepsilon_0^a \text{ and } U_0^{\text{DDC}}(A_1) = \delta \mathbb{E} [\max\{v(a) + \varepsilon_1^a, v(b) + \varepsilon_1^b\}] + \varepsilon_0^{A_1}.$$

Section 5 shows that shocks to actions can lead to counterintuitive behavior, such as a negative option value. Additionally, they can result in biased parameter estimates: In the present example, the maximum likelihood estimate of the discount factor under i.i.d. DDC is exaggerated relative to BEU. ▲

2 Static vs. Dynamic Random Utility

For any set Y , denote by $\mathcal{K}(Y)$ the set of all nonempty finite subsets of Y and by $\Delta(Y)$ the set of all simple (i.e., finite support) lotteries on Y ; henceforth, all references to lotteries are to simple lotteries. Whenever Y is a separable metric space, we endow $\Delta(Y)$ with the induced Prokhorov metric and $\mathcal{K}(Y)$ with the Hausdorff metric. Let \mathbb{R}^Y denote the set of vNM utility indices over Y , which is endowed with the product topology and its induced Borel sigma-algebra. For any $U, U' \in \mathbb{R}^Y$, write $U \approx U'$ if U and U' represent the same preference on $\Delta(Y)$. For any finite set of lotteries $A \in \mathcal{K}(\Delta(Y))$, let $M(A, U) := \operatorname{argmax}_{p \in A} U(p)$ denote the set of lotteries in A that maximize U , where $U(p) := \sum_{y \in \operatorname{supp}(p)} U(y)p(y)$ denotes the expected utility of any $p \in \Delta(Y)$. For any $A, B \in \mathcal{K}(\Delta(Y))$ and $\alpha \in [0, 1]$, define the α -mixture of A and B by $\alpha A + (1 - \alpha)B := \{\alpha p + (1 - \alpha)q : p \in A, q \in B\} \in \mathcal{K}(\Delta(Y))$.

2.1 Static Random Utility

We first briefly review the static model of random expected utility that will serve as the building block of our dynamic representation at each history. The model is based on [Gul and Pesendorfer \(2006\)](#), but allows for an infinite outcome space; this extension is necessary for our purposes, because in the dynamic setting the period- t outcome space X_t , consisting of all pairs of current consumptions and continuation menus, will be infinite in all but the final period. In [Section 2.2.3](#), we interpret the stochastic choice data that this model gives rise to in terms of a large population of heterogeneous individuals whose identities are unknown to the analyst, but by convention, we express the model and axioms in terms of the behavior of any one of these individuals, referred to as “the agent.”

2.1.1 Agent’s problem

Let X be an arbitrary separable metric space of outcomes. The agent makes choices from menus, which are finite sets of lotteries over X ; the set of all menus is $\mathcal{A} := \mathcal{K}(\Delta(X))$. Denote a typical menu by A and a typical lottery by p . Let $(\Omega, \mathcal{F}^*, \mu)$ be a finitely-additive probability space. In each state of the world, the agent’s choices maximize her expected utility subject to her private information. Her payoff-relevant private information is captured by a sigma-algebra $\mathcal{F} \subseteq \mathcal{F}^*$ and an \mathcal{F} -measurable random vNM utility index $U : \Omega \rightarrow \mathbb{R}^X$. In case of indifference, ties are broken by a random vNM index $W : \Omega \rightarrow \mathbb{R}^X$, which is measurable with respect to \mathcal{F}^* . Thus, when faced with menu A , the agent chooses lottery p in state ω if and only if p maximizes $U(\omega)$ in A and, in case of ties, additionally maximizes $W(\omega)$ among the $U(\omega)$ -maximizers. The event in which the agent chooses p from A is $C(p, A) := \{\omega \in \Omega : p \in M(M(A, U(\omega)), W(\omega))\}$.

For tractability, we follow [Ahn and Sarver \(2013\)](#) in assuming that the agent’s payoff-relevant private information (\mathcal{F}, U) is *simple*, i.e., (i) \mathcal{F} is generated by a finite partition such that $\mu(\mathcal{F}(\omega)) > 0$ for every $\omega \in \Omega$, where $\mathcal{F}(\omega)$ denotes the cell of the partition that contains ω ; and (ii) each $U(\omega)$ is nonconstant and $U(\omega) \not\approx U(\omega')$ whenever $\mathcal{F}(\omega) \neq \mathcal{F}(\omega')$. Moreover, the tie-breaker W is *proper*,⁸ ensuring that under W ties occur with probability 0 in each menu; that is, $\mu(\{\omega \in \Omega : |M(A, W(\omega))| = 1\}) = 1$ for all $A \in \mathcal{A}$.

2.1.2 Analyst’s problem

The analyst does not observe the agent’s private information and thus cannot condition on events in \mathcal{F} . Because of this informational asymmetry, the agent’s choices appear stochastic to the analyst.⁹ His observations are summarized by a *stochastic choice rule* on \mathcal{A} , i.e., a map $\rho : \mathcal{A} \rightarrow \Delta(\Delta(X))$ such that $\sum_{p \in A} \rho(p, A) = 1$ for all $A \in \mathcal{A}$. Here $\rho(p, A)$ denotes the probability with which the agent chooses lottery p when faced with menu A . If the agent behaves as in the previous section, then the event that the agent chooses p from A is $C(p, A)$. Thus, the analyst’s observations are consistent with the previous section if $\rho(p, A) = \mu(C(p, A))$ for all p and A .

Definition 1. A *static random expected utility (REU)* representation of the stochastic choice rule ρ is a tuple $(\Omega, \mathcal{F}^*, \mu, \mathcal{F}, U, W)$ such that $(\Omega, \mathcal{F}^*, \mu)$ is a finitely-additive probability space, the sigma-algebra $\mathcal{F} \subseteq \mathcal{F}^*$ and the \mathcal{F} -measurable utility $U : \Omega \rightarrow \mathbb{R}^X$ are simple, the \mathcal{F}^* -measurable tiebreaker $W : \Omega \rightarrow \mathbb{R}^X$ is proper, and $\rho(p, A) = \mu(C(p, A))$ for all p and A .

⁸This property is sometimes called “regular” in the literature; we use the term “proper” to avoid confusion with the Regularity axiom (Axiom 0 (i)) below.

⁹If the analyst observed the true state, choices would appear deterministic and could be summarized by a vNM preference \succsim_ω .

2.1.3 Characterization

For finite outcome spaces X , static REU representations have been characterized by Gul and Pesendorfer (2006) and Ahn and Sarver (2013). As a preliminary technical contribution, we extend their characterization to simple lotteries over *arbitrary* separable metric spaces X . The first four conditions of the following axiom are the same as in Gul and Pesendorfer (2006). The fifth condition is a slight modification of the finiteness condition in Ahn and Sarver (2013).

Axiom 0. (Random Expected Utility)

- (i). *Regularity*: If $A \subseteq A'$, then for all $p \in A$, $\rho(p; A) \geq \rho(p; A')$.
- (ii). *Linearity*: For any $A, p \in A, \lambda \in (0, 1)$, and q , $\rho(p; A) = \rho(\lambda p + (1 - \lambda)q; \lambda A + (1 - \lambda)\{q\})$.
- (iii). *Extremeness*: For any A , $\rho(\text{ext}A; A) = 1$.¹⁰
- (iv). *Mixture Continuity*: $\rho(\cdot; \alpha A + (1 - \alpha)A')$ is continuous in α for all A, A' .
- (v). *Finiteness*: There is $K > 0$ such that for all A , there is $B \subseteq A$ with $|B| \leq K$ such that for every $p \in A \setminus B$, there are sequences $p^n \rightarrow^m p$ and $B^n \rightarrow^m B$ with $\rho(p^n; \{p^n\} \cup B^n) = 0$ for all n .

For condition (iv), $\alpha \mapsto \rho(\cdot; \alpha A + (1 - \alpha)A')$ is viewed as a map from $[0, 1]$ to $\Delta(\Delta(X))$, where $\Delta(\Delta(X))$ is endowed with the topology of weak convergence induced by the Prokhorov metric on $\Delta(X)$. For condition (v), *convergence in mixture*, denoted \rightarrow^m , on $\Delta(X)$ and \mathcal{A} is defined as follows: For any $p \in \Delta(X)$ and sequence $\{p^n\}_{n \in \mathbb{N}} \subseteq \Delta(X)$, we write $p^n \rightarrow^m p$ if there exists $q \in \Delta(X)$ and a sequence $\{\alpha_n\}_{n \in \mathbb{N}}$ with $\alpha_n \rightarrow 0$ such that $p^n = \alpha_n q + (1 - \alpha_n)p$ for all n . Similarly, for any sequence $\{B^n\}_{n \in \mathbb{N}} \subseteq \mathcal{A}$, we write $B^n \rightarrow^m p$ if there exists $B \in \mathcal{A}$ and a sequence $\{\alpha_n\}_{n \in \mathbb{N}}$ with $\alpha_n \rightarrow 0$ such that $B^n = \alpha_n B + (1 - \alpha_n)\{p\}$ for all n . Finally, for any $A \in \mathcal{A}$ and sequence $(A^n)_{n \in \mathbb{N}} \subseteq \mathcal{A}$, we write $A^n \rightarrow^m A$ if for each $p \in A$, there is a sequence $\{B_p^n\}_{n \in \mathbb{N}} \subseteq \mathcal{A}$ such that $B_p^n \rightarrow^m p$ and $A^n = \cup_{p \in A} B_p^n$ for all n .

Theorem 0. The stochastic choice rule ρ on \mathcal{A} satisfies Axiom 0 if and only if ρ admits an REU representation.

Proof. See Supplementary Appendix F. ■

2.2 Dynamic Random Utility

Motivated by the examples in Section 1.3, in what follows, we set up and characterize a general model of dynamic random utility.

¹⁰Here $\text{ext}A$ denotes the set of extreme points of A .

2.2.1 Agent’s Problem

The agent faces a decision tree, as defined by [Kreps and Porteus \(1978\)](#). There are finitely many periods $t = 0, 1, \dots, T$. There is a finite set Z of instantaneous consumptions. Each period t , the agent chooses from a period- t menu, which is a finite set of lotteries over the period- t outcome space X_t . The spaces X_t are defined recursively. The final period outcome space $X_T := Z$ is just the space of instantaneous consumptions; the set of all period- T menus is $\mathcal{A}_T := \mathcal{K}(\Delta(X_T))$. In all earlier periods $t \leq T - 1$, the outcome space $X_t := Z \times \mathcal{A}_{t+1}$ consists of all pairs of current period consumptions and next period continuation menus; the set of period- t menus is $\mathcal{A}_t := \mathcal{K}(\Delta(X_t))$.¹¹ Denote a typical period- t lottery by $p_t \in \Delta(X_t)$ and a typical menu by $A_t \in \mathcal{A}_t$. The agent’s choice of $p_t \in A_t$ determines both her instantaneous consumption z_t and the menu A_{t+1} from which she will choose next period; let $p_t^Z \in \Delta(Z)$ and $p_t^A \in \Delta(\mathcal{A}_{t+1})$ denote the respective marginal distributions.

As in the static model, let $(\Omega, \mathcal{F}^*, \mu)$ be a finitely-additive probability space. Under *dynamic random expected utility (DREU)*, in each state of the world and in each period, the agent’s choices maximize her expected utility subject to her dynamically evolving private information. The agent’s payoff-relevant private information is captured by a filtration $(\mathcal{F}_t)_{0 \leq t \leq T} \subseteq \mathcal{F}^*$ and an \mathcal{F}_t -adapted process of random vNM utility indices $U_t : \Omega \rightarrow \mathbb{R}^{X_t}$ over X_t . This allows for arbitrary serial correlation of utilities, but does not allow the utility process to depend on past consumption; [Section 7.2](#) discusses how to relax the latter restriction. In case of indifference, ties at each t are broken by a random \mathcal{F}^* -measurable vNM utility index $W_t : \Omega \rightarrow X_t$, where we impose dynamic analogs of simplicity and properness that we define at the end of this section. Thus, as before, when faced with menu A_t in period t , the agent chooses lottery p_t in the event $C(p_t, A_t) := \{\omega \in \Omega : p_t \in M(A_t, U_t(\omega)), W_t(\omega)\}$.

DREU is a very general model because it imposes no particular structure on the family (U_t) . This is the most parsimonious setting in which to isolate the behavioral implications of serially correlated private information. DREU could also accommodate various behavioral effects, such as temptation or certain forms of “mistakes” (e.g., [Ke, 2018](#)), which in the static setting are indistinguishable from random utility maximization. However, the following important special case rules out these possibilities.

Bayesian evolving utility (BEU) captures a dynamically sophisticated agent who correctly takes into account the evolution of her future preferences. There is an \mathcal{F}_t -adapted process of random felicity functions $u_t : \Omega \rightarrow \mathbb{R}^Z$ over instantaneous consumptions and an \mathcal{F}_t -adapted process of discount factors $\delta_t : \Omega \rightarrow \mathbb{R}_{++}$ such that $U_T = u_T$ and U_t for $t \leq T - 1$ is given by

¹¹A small technical difference from [Kreps and Porteus \(1978\)](#) is that they use Borel instead of simple lotteries and compact instead of finite menus, but as in their setting we can verify recursively that each X_t is a separable metric space under the appropriate topologies (see [Lemma E.1](#)).

the Bellman equation

$$U_t(z_t, A_{t+1}) = u_t(z_t) + \delta_t \mathbb{E} \left[\max_{p_{t+1} \in A_{t+1}} U_{t+1}(p_{t+1}) | \mathcal{F}_t \right]. \quad (1)$$

In (1) the process of discount factors is not identified. An important special case of BEU where the process δ_t is identified is *Bayesian evolving beliefs (BEB)*.¹² This captures the setting, discussed in Example 1, where the agent has a fixed but unknown felicity about which she learns over time. Formally, there is an \mathcal{F}^* -measurable random felicity $\tilde{u} : \Omega \rightarrow \mathbb{R}^Z$ such that for all t ¹³

$$u_t = \mathbb{E}[\tilde{u} | \mathcal{F}_t]. \quad (2)$$

For all three models, we impose the following dynamic analogs of simplicity and properness. The pair $(\mathcal{F}_t, U_t)_{0 \leq t \leq T}$ is *simple*, i.e., (i) each \mathcal{F}_t is generated by a finite partition such that $\mu(\mathcal{F}_t(\omega)) > 0$ for every $\omega \in \Omega$, where $\mathcal{F}_t(\omega)$ again denotes the cell of the partition that contains ω ; and (ii) each $U_t(\omega)$ is nonconstant, and $U_t(\omega) \not\approx U_t(\omega')$ whenever $\mathcal{F}_t(\omega) \neq \mathcal{F}_t(\omega')$ and $\mathcal{F}_{t-1}(\omega) = \mathcal{F}_{t-1}(\omega')$.¹⁴ The tiebreakers $(W_t)_{0 \leq t \leq T}$ are *proper*, i.e., (i) $\mu(\{\omega \in \Omega : |M(A_t, W_t(\omega))| = 1\}) = 1$ for all $A_t \in \mathcal{A}_t$; (ii) conditional on $\mathcal{F}_T(\omega)$, W_0, \dots, W_T are independent; and (iii) $\mu(W_t \in B_t | \mathcal{F}_T(\omega)) = \mu(W_t \in B_t | \mathcal{F}_t(\omega))$ for all t and measurable B_t .¹⁵

2.2.2 Analyst's Problem

As in the static setting, the agent's choices in each period t appear stochastic to the analyst, because he does not have access to the agent's private information. The novel feature of the dynamic setting is that the analyst can observe the agent's past choices. With serially correlated utilities, these choices convey some information about the payoff-relevant private information \mathcal{F}_t , so that the agent's behavior additionally appears *history dependent* to the analyst.

This is captured by a *dynamic stochastic choice rule* ρ , which for any period t and history of past choices summarizes the observed choice frequencies from any menu A_t that can arise after this history. We define choice frequencies and histories recursively. Choice frequencies in period 0 are given by a (static) stochastic choice rule $\rho_0 : \mathcal{A}_0 \rightarrow \Delta(\Delta(X_0))$ on \mathcal{A}_0 ; thus, $\rho_0(p_0; A_0)$ denotes the probability with which the agent chooses lottery p_0 when faced with

¹²We allow for the possibility that discount factors are stochastic and/or evolving, but it is straightforward to characterize the case of a fixed discount factor $\delta \in \mathbb{R}_{++}$. See the discussion following Theorem 3.

¹³BEB is a model of passive learning, because the agent's choices do not affect her filtration \mathcal{F}_t . A consumption-dependent extension of BEB (see Section 7.2) can accommodate active learning/experimentation, where each period the agent obtains additional information from her consumption z_t .

¹⁴For $t = 0$, we let $\mathcal{F}_{t-1}(\omega) := \Omega$ for all ω .

¹⁵(ii) rules out additional serial correlation of tiebreakers, over and above the serial correlation inherent in the agent's payoff-relevant private information $\mathcal{F}_T(\omega)$. (iii) ensures that to the extent that period- t tie breaking relies on payoff-relevant private information, it can rely only on the information $\mathcal{F}_t(\omega)$ available at t .

menu A_0 . The choices that occur with strictly positive probability under ρ_0 define the set of all period 0 *histories* $\mathcal{H}_0 := \{(A_0, p_0) : \rho_0(p_0, A_0) > 0\}$. For any history $h^0 = (A_0, p_0) \in \mathcal{H}_0$, let $\mathcal{A}_1(h^0) := \text{supp } p_0^A$ denote the set of period 1 menus that follow h^0 with positive probability.

For $t \geq 1$ the objects \mathcal{H}_t and $\mathcal{A}_{t+1}(h^t)$ are defined recursively. For any history $h^{t-1} \in \mathcal{H}_{t-1}$, choice frequencies following h^{t-1} are given by a stochastic choice rule $\rho_t(\cdot | h^{t-1}) : \mathcal{A}_t(h^{t-1}) \rightarrow \Delta(\Delta(X_t))$ on the set $\mathcal{A}_t(h^{t-1})$ of period t menus that follow h^{t-1} with positive probability; thus, $\rho_t(p_t; A_t | h^{t-1})$ denotes the probability with which the agent chooses p_t when faced with menu A_t after history h^{t-1} . The set of period- t histories is $\mathcal{H}_t := \{(h^{t-1}, A_t, p_t) : h^{t-1} \in \mathcal{H}_{t-1} \text{ and } A_t \in \mathcal{A}_t(h^{t-1}) \text{ and } \rho_t(p_t; A_t | h^{t-1}) > 0\}$; this contains all sequences $(A_0, p_0, \dots, A_t, p_t)$ of choices up to time t that arise with positive probability. Finally, for each $t \leq T - 1$, the set of period $t + 1$ menus that follow history $h^t = (h^{t-1}, A_t, p_t)$ with positive probability is $\mathcal{A}_{t+1}(h^t) := \text{supp } p_t^A$ and the set of period- t histories that lead to A_{t+1} with positive probability is $\mathcal{H}_t(A_{t+1}) := \{h^t \in \mathcal{H}_t : A_{t+1} \in \mathcal{A}_{t+1}(h^t)\}$.

Two features of the primitive are worth noting: First, for each $t \geq 1$ and history $h^{t-1} \in \mathcal{H}_{t-1}$, the stochastic choice rule $\rho_t(\cdot | h^{t-1})$ is defined only on the subset $\mathcal{A}_t(h^{t-1}) \subseteq \mathcal{A}_t$ of period t menus that arise with positive probability after h^{t-1} —typically very few menus. This reflects a key property of the decision-tree formulation that we term *limited observability*, whereby histories of choices also encode all possible future menus that the agent will face, as illustrated in Example 2. Nevertheless, Section 3.2 will show that under DREU the analyst can extrapolate from $\rho_t(\cdot | h^{t-1})$ to a well-defined stochastic choice rule on the whole of \mathcal{A}_t . Second, histories only summarize the agent’s past choices of p_k from A_k and do not keep track of realized consumptions $z_k \in \text{supp } p_k^Z$. This is without loss in the current model where utilities are not affected by past consumption, but Section 7.2 discusses a generalization of our model that relaxes this assumption.

Under DREU, the private information revealed to the analyst by history $h^{t-1} = (A_0, p_0, \dots, A_{t-1}, p_{t-1})$ is given by the event $C(h^{t-1}) := \bigcap_{k=0}^{t-1} C(p_k, A_k)$.¹⁶ Thus, the analyst’s observations are consistent with DREU if the probability with which the agent chooses p_t from A_t following history h^{t-1} is equal to the conditional probability $\mu[C(p_t, A_t) | C(h^{t-1})]$ of the event $C(p_t, A_t)$ given $C(h^{t-1})$.

The following definition summarizes the dynamic model:

Definition 2. A *dynamic random expected utility (DREU)* representation of the dynamic stochastic choice rule ρ is a tuple $(\Omega, \mathcal{F}^*, \mu, (\mathcal{F}_t, U_t, W_t)_{0 \leq t \leq T})$ such that $(\Omega, \mathcal{F}^*, \mu)$ is a finitely-additive probability space, the filtration $(\mathcal{F}_t) \subseteq \mathcal{F}^*$ and the \mathcal{F}_t -adapted utility process $U_t : \Omega \rightarrow \mathbb{R}^{X_t}$ are simple, the \mathcal{F}^* -measurable tiebreaking process $W_t : \Omega \rightarrow \mathbb{R}^{X_t}$ is proper,

¹⁶Note that $C(h^{t-1})$ does not keep track of the random realizations of menus $A_k \in \text{supp } p_k^A$ along the sequence h^{t-1} , as this exogenous randomness does not reveal any information about the agent’s private information.

and for all $p_t \in A_t$ and $h^{t-1} \in \mathcal{H}_{t-1}(A_t)$,

$$\rho_t(p_t; A_t | h^{t-1}) = \mu [C(p_t, A_t) | C(h^{t-1})], \quad (3)$$

where for $t = 0$, we abuse notation by letting $C(h^{t-1}) := \Omega$ and $\rho_0(p_0; A_0 | h^{-1}) := \rho_0(p_0; A_0)$.

A *Bayesian evolving utility (BEU)* representation is a DREU representation along with \mathcal{F}_t -adapted processes of felicities $u_t : \Omega \rightarrow \mathbb{R}^Z$ and discount factors $\delta_t : \Omega \rightarrow \mathbb{R}_{++}$ such that (1) holds. A *Bayesian evolving beliefs (BEB)* representation is a BEU representation along with an \mathcal{F}^* -measurable felicity $\tilde{u} : \Omega \rightarrow \mathbb{R}^Z$ such that (2) holds.

2.2.3 Discussion

Lotteries as choice objects. In addition to allowing us to model choice behavior under risk, including lotteries in the domain of choice simplifies our analysis, as it allows us to rely on the static framework of [Gul and Pesendorfer \(2006\)](#) instead of the more complicated one of [Falmagne \(1978\)](#). Lotteries play a similar technical role in the original work of [Kreps and Porteus \(1978\)](#), by letting them rely on the vNM framework.¹⁷ From a conceptual point of view, we will see in [Section 3.2](#) that lotteries are crucial in overcoming the aforementioned limited observability problem and we illustrate the availability of lotteries for this purpose with examples from experimental and empirical work.

Interpretation of data. We interpret the dynamic stochastic choice data ρ as the analyst's observation of a large population of agents that solve each decision tree once; agents have heterogeneous and evolving utilities that are realized independently according to the model in [Section 2.2](#) and the analyst does not observe agents' identities (only their choice histories). This interpretation resembles available data in empirical analysis. However, (analogous to the static setting) the results do not rule out an alternative interpretation, whereby the analyst observes a single agent solve each decision tree repeatedly.¹⁸ In either case, ρ captures the limiting choice frequencies as the population size/number of observations tends to infinity. Abstracting from the sampling error in this manner is also typical in the econometric analysis of identification. In any application, the data set will of course be finite. However, studying behavior on the full domain is an important step in uncovering all the assumptions that are behind the model; moreover, statistical tests are often directly inspired by axioms.¹⁹

Dynamic stochastic choice vs. ex ante preference. In our framework, the analyst

¹⁷Likewise, the ambiguity aversion literature extensively relies on the [Anscombe and Aumann \(1963\)](#) framework rather than the more complicated one of [Savage \(1972\)](#); the notable exceptions include [Gilboa \(1987\)](#) and [Epstein \(1999\)](#). Similarly, the menu-preference literature uses lotteries (e.g. [Dekel, Lipman, and Rustichini, 2001](#)) to improve upon the uniqueness and comparative statics results of [Kreps \(1979\)](#).

¹⁸Here, the agent's utilities are assumed to evolve according to the same process U_t at each observation.

¹⁹For example, [Hausman and McFadden \(1984\)](#) develop a test of the IIA axiom that characterizes the logit model. Likewise, [Kitamura and Stoye \(2018\)](#) develop axiom-based tests of the static random utility model.

observes the distribution of choices at each node of each decision tree; as we pointed out, the randomness in choice comes from an informational asymmetry between agents and the analyst in each period. By contrast, a widespread approach in the existing dynamic decision theory literature (e.g., Gul and Pesendorfer, 2004; Krishna and Sadowski, 2014) is to only study a deterministic preference over decision trees at a hypothetical ex ante stage that features no informational asymmetry²⁰ or abstracts away from other forces (e.g., temptation) that the agent anticipates to affect her choices in actual decision trees.²¹ Compared with this literature, our approach does not require such a hypothetical stage, and thus the primitive is closer to actual data economists can observe. Moreover, considering choice behavior in each period, not just at the beginning of time, allows us to study phenomena such as history dependence and choice persistence and to test whether the agent’s expectations are correct.

Role of axioms. In addition to their usual positive and normative role, we view our axioms as serving an equally important purpose as conceptual tools that elucidate key properties of any dynamic random utility model and facilitate comparisons between different versions of the model. For example, our axioms in Section 3.1 clarify the nature of history dependence that can arise under any dynamic random expected utility model; our axioms in Sections 4.2 and 6.2 identify the additional behavioral content of Bayesian evolving beliefs relative to Bayesian evolving utility; and our comparison of BEU and i.i.d. DDC in Section 5 draws on the axioms to uncover that the two make opposite predictions about option value.

3 Characterization of DREU

DREU is characterized by four axioms, which we present in the following subsections. First, we present two history independence axioms that capture the key new implications of the dynamic model relative to the static one. Building on this, the next subsection shows how the analyst can extrapolate from each $\rho_t(\cdot|h^{t-1})$ to an extended choice rule on the whole of \mathcal{A}_t , thus overcoming the limited observability problem. The final subsection then imposes the static REU conditions as well as a technical history continuity axiom on this extended choice rule.

For simplicity of exposition, we present our characterization in the two-period setting ($T = 1$); the generalization to an arbitrary finite horizon is straightforward and is provided in Appendix B.1.

²⁰Ahn and Sarver (2013) study a two-period model with a deterministic menu preference in the first period and random choice from menus in the second period. Here too there is no informational asymmetry in the first period.

²¹In the context of temptation, one exception is Noor (2011), but his is a stationary environment with no informational asymmetry and the analyst observes deterministic choices at each node of the decision tree.

3.1 History Independence Axioms

Our first two axioms identify two cases in which histories h^0 and g^0 reveal the same information to the analyst. Capturing the fact that history dependence arises in DREU only through the private information revealed by past choices, the axioms require that period-1 choice behavior be the same after two such histories.

First, consider two histories $h^0 = (A_0, p_0)$ and $g^0 = (B_0, p_0)$ that differ solely in that $A_0 \subseteq B_0$ is a contraction of B_0 , and suppose that both histories arise with the same probability $\rho_0(p_0; A_0) = \rho_0(p_0; B_0)$. Axiom 1 requires period-1 choice behavior to be the same after h^0 and g^0 .

Axiom 1 (Contraction History Independence). If $(A_0, p_0) \in \mathcal{H}_0(A_1)$ and $B_0 \supseteq A_0$ with $\rho_0(p_0; A_0) = \rho_0(p_0; B_0)$, then $\rho_1(\cdot; A_1 | A_0, p_0) = \rho_1(\cdot; A_1 | B_0, p_0)$.

To see the idea, note that in general, the event that p_0 is the best element of menu B_0 is a subset of the event that p_0 is the best element of the smaller menu $A_0 \subseteq B_0$; thus, observing $g^0 = (B_0, p_0)$ may reveal more information about the agent's possible period-0 preferences than $h^0 = (A_0, p_0)$. However, since we additionally know that $\rho_0(p_0; A_0) = \rho_0(p_0; B_0)$, the event that p_0 is best in A_0 but not in B_0 must have probability 0; in other words, we must put zero probability on any preference that selects p_0 from A_0 but not from B_0 . Given this, h^0 and g^0 reveal the same information, and hence call for the same predictions for period-1 choices. The following example illustrates Axiom 1 in a simple setting where agents only choose instantaneous consumption in each period and today's choice does not affect tomorrow's menu.²²

Example 4. Consider again the brand choice data from Example 1. Suppose the left and right panel of Figure 1 respectively represent purchasing data at two stores, A and B . Both stores typically carry two brands of milk, non-organic (x) and organic (y), but in week 0, store B exceptionally offers an additional organic brand z . The week-0 purchasing shares at each store are as in Figure 1. In particular, the share of customers purchasing the non-organic brand x in week 0 is the *same* (80%) at both stores. Assume each store has a stable set of weekly customers whose stochastic process of preferences is identical at both stores.²³

If in week 1 both stores carry only x and y , then Contraction History Independence implies that the week-1 choice frequencies among customers who bought x in week 0 must be the same at both stores. Indeed, consider any customers Alice of store A and Barbara of store B who both buy brand x in week 0. Then we have the same information about Alice and Barbara. Since at store A only x and y were available in week 0, the possible week-0 preferences of Alice are $x \succ y \succ z$ or $x \succ z \succ y$ or $z \succ x \succ y$. By contrast, since store B stocked all three brands,

²²Section 6 studies this domain of “atemporal consumption problems” in more detail.

²³For simplicity, we assume in the following that all preferences are strict.

Barbara’s possible preferences are $x \succ y \succ z$ or $x \succ z \succ y$. However, since we additionally know that the week-0 demand share of brand x was the same at both stores, $\rho_0(x; \{x, y, z\}) = \rho_0(x; \{x, y\}) = 0.8$, we can conclude that *no* customers had the ranking $z \succ x \succ y$ in week 0. Therefore, the analyst’s prediction is the same, since the stochastic process that governs the transition from week-0 to week-1 preferences is the same for Barbara and Alice and in both cases the analyst conditions on exactly the same week-0 event $\{x \succ y \succ z, x \succ z \succ y\}$. \blacktriangle

This is similar to the idea that in the static setting, Regularity (Axiom 0 (i)) rules out certain “irrational” behavior such as the attraction effect (e.g., Huber, Payne, and Puto, 1982), where the mere addition of some unchosen decoy option affects the agent’s choice probabilities over existing options. Likewise, Contraction History Independence rules out certain *dynamically* irrational choice patterns such as the “mere exposure effect” (e.g., Zajonc, 2001), where an agent’s choices today are influenced by the mere availability of irrelevant options in the *past*.²⁴ For instance, in Example 4, the axiom rules out the possibility that Barbara’s choices in week 1 are affected by merely seeing (but not buying) brand z in week 0.

Contraction History Independence only concerns histories h^0 and g^0 that share the same past choice p_0 . Our second history independence axiom imposes discipline across certain histories that feature different choices. This axiom takes into account the fact that the agent is an expected utility maximizer. Under expected utility maximization, choosing p_0 from A_0 reveals the same information about the agent’s utility as choosing $\lambda p_0 + (1 - \lambda)q_0$ from $\lambda A_0 + (1 - \lambda)\{q_0\}$. Thus, period-1 choice behavior following history $h^0 = (A_0, p_0)$ or history $g^0 = (\lambda A_0 + (1 - \lambda)\{q_0\}, \lambda p_0 + (1 - \lambda)q_0)$ should be the same. For instance, in the school choice example (Example 2), parents who in Figure 2 (left) chose school 1 should make the same choices from the resulting period-1 menu $\{H, P, S\}$ as parents who in Figure 2 (right) chose the lottery $\lambda(\text{school 1}) + (1 - \lambda)(\text{school 2})$ and were allocated to school 1.

More generally, for any menu B_0 , if we know that the agent chose *some* option of the form $\lambda p_0 + (1 - \lambda)q_0$ from $\lambda A_0 + (1 - \lambda)B_0$ but we do not know what q_0 was, this again reveals the same information as choosing p_0 from A_0 . Thus, conditioning on history h^0 or on the *set* of histories $G^0 = \{\lambda h^0 + (1 - \lambda)(B_0, q_0) : q_0 \in B_0\}$ should again yield the same predictions for period-1 choice behavior, where $\lambda h^0 + (1 - \lambda)(B_0, q_0)$ is shorthand for $(\lambda A_0 + (1 - \lambda)B_0, \lambda p_0 + (1 - \lambda)q_0)$.²⁵

This is the content of Axiom 2. To state this formally, define the choice distribution from A_1 following any set of histories $G^0 \subseteq \mathcal{H}_0(A_1)$,

$$\rho_1(\cdot; A_1 | G^0) := \sum_{g^0 \in G^0} \rho_1(\cdot; A_1 | g^0) \frac{\rho_0(g^0)}{\sum_{f^0 \in G^0} \rho_0(f^0)},$$

²⁴Cerigioni (2017) incorporates the exposure effect into a Luce-style model in a dynamic setting.

²⁵The proof sketch of Theorem 1 in Section 3.4 illustrates the role played by allowing for sets of histories G^0 , rather than only singleton histories $g^0 = \lambda h^0 + (1 - \lambda)(\{q_0\}, q_0)$ in Axiom 2.

to be the weighted average of all choice distributions $\rho_1(\cdot; A_1|g^0)$ following individual histories in G^0 , where each history $g^0 = (\hat{A}_0, \hat{p}_0)$ is weighted by the probability that it arises, $\rho_0(g^0) := \rho_0(\hat{p}_0; \hat{A}_0)$.

Axiom 2 (Linear History Independence). If $h^0 \in \mathcal{H}_0(A_1)$ and $G^0 = \{\lambda h^0 + (1 - \lambda)(B_0, q_0) : q_0 \in B_0\} \subseteq \mathcal{H}_0$ for some $\lambda \in (0, 1]$, then $\rho_1(\cdot; A_1|h^0) = \rho_1(\cdot; A_1|G^0)$.

A number of recent experimental studies feature the following type of setting that allows for a simple test of Axiom 2: In period 0, subjects are presented with the choice between (i) some period-1 menu B_1 and (ii) a lottery that yields some other period-1 menu A_1 with probability λ and menu B_1 with probability $1 - \lambda$; in period 1, subjects make choices from their realized menus.²⁶ Here Linear History Independence implies that period-1 choices (from A_1 or B_1) among subjects who choose (ii) in period 0 should be independent of the particular value of $\lambda \in (0, 1]$; this can be tested by exogenously varying this randomization probability.

3.2 Limited Observability

Recall that unlike the static setting, where the analyst observes choices from all possible menus, the dynamic setting presents a limited observability problem: At each history h^0 , $\rho_1(\cdot|h^0)$ is only defined on the set $\mathcal{A}_1(h^0)$ of menus that occur with positive probability after h^0 —typically very few menus. For the rest of the paper, it is key to overcome this problem: Otherwise we do not have enough data to verify whether observed choices at history h^0 are consistent with random utility maximization or to identify whether the agent’s utility process belongs to the Bayesian evolving utility class or the more specific evolving beliefs class.

The inclusion of lotteries among the agent’s choice objects allows us to do so. In particular, Linear History Independence provides a formal justification for the “linear extrapolation” procedure illustrated in the school choice example (Example 2). Consider any menu A_1 (e.g., the two-option menu $\{H, P\}$ in the example) and some history $h^0 = (A_0, p_0)$ that does not lead to A_1 (e.g., choosing school 1 from menu $\{\text{school 1}, \text{school 2}\}$ in the left-hand tree in Figure 2). To define the agent’s counterfactual choice distribution from A_1 following h^0 , we extrapolate from a situation where the agent knows that no matter which option in A_0 she chooses, with some fixed probability another option q_0 that *does* lead to menu A_1 will be implemented instead.

More precisely, we pick a lottery q_0 such that $A_1 \in \text{supp } q_0^A$ and replace menu A_0 with $\lambda A_0 + (1 - \lambda)\{q_0\}$. This corresponds to the right-hand tree in Figure 2, where the choice between school 1 and school 2 is replaced with the choice between the *lottery* $\lambda(\text{school 1}) + (1 - \lambda)(\text{school 2})$

²⁶E.g., Toussaert’s (2018) recent experiment on temptation and self-control uses a similar design to differentiate between so-called random Strotz agents and Gul and Pesendorfer (2001) agents. Related uses of randomization over menus in lab experiments include Augenblick, Niederle, and Sprenger (2015); Dean and McNeill (2016). To avoid certainty effects, these experiments typically do not feature any degenerate lotteries as in (i), but we abstract away from this for expositional simplicity.

and school 2.²⁷ As discussed preceding Linear History Independence, under expected utility maximization, choosing p_0 from A_0 reveals the same information about the agent as choosing $\lambda p_0 + (1 - \lambda)q_0$ from $\lambda A_0 + (1 - \lambda)\{q_0\}$. This motivates defining choice behavior from A_1 following history $h^0 = (A_0, p_0)$ by extrapolating from choices following history $g^0 = \lambda h^0 + (1 - \lambda)(\{q_0\}, q_0)$:

Definition 3. For any $A_1 \in \mathcal{A}_1$, and $h^0 \in \mathcal{H}_0$, define

$$\rho_1^{h^0}(\cdot; A_1) := \rho_1(\cdot; A_1 | \lambda h^0 + (1 - \lambda)(\{q_0\}, q_0)) \quad (4)$$

for some $\lambda \in (0, 1]$ and q_0 with $A_1 \in \text{supp } q_0^A$.

Linear History Independence justifies Definition 3, as it ensures that the extended choice rule $\rho_1^{h^0}(\cdot; A_1)$ is well-defined: Lemma E.4 shows that the RHS of (4) does not depend on the specific choice of λ and q_0 ; moreover, $\rho_1^{h^0}(\cdot; A_1)$ coincides with $\rho_1(\cdot; A_1 | h^0)$ whenever $h^0 \in \mathcal{H}_0(A_1)$. In the following, we do not distinguish between the extended and nonextended version of ρ_1 and use $\rho_1(\cdot; A_1 | h^0)$ to denote both.

As Example 2 illustrates in the context of school choice, random assignment is prevalent in many real-world economic environments and is an important tool to obtain quasi-experimental variation in the empirical literature. While this literature typically leverages such random variation to identify the causal effect of current choices on next-period *outcomes* (e.g., test scores in the case of school choice), Definition 3 suggests exploiting it to make counterfactual inferences about next-period *choices*. Even more readily, lotteries over next-period choice problems can be generated in the laboratory; as discussed following Axiom 2, a growing literature in experimental economics features this type of randomization, and one purpose is precisely to perform extrapolation procedures akin to Definition 3.

3.3 History-Dependent REU and History Continuity Axioms

For each h^0 , the extended choice distribution $\rho_1(\cdot | h^0)$ from Definition 3 is a stochastic choice rule on the whole of \mathcal{A}_1 . The next axiom imposes the standard static REU conditions from Axiom 0 on ρ_0 as well as on each $\rho_1(\cdot | h^0)$.²⁸ Note that conditioning ρ_1 on period-0 histories is essential; without controlling for past choices, period-1 choice behavior will in general violate the REU axioms, as illustrated in Example 2.

²⁷Note that by definition, menu $\{\lambda(\text{school 1}) + (1 - \lambda)(\text{school 2}), \text{school 2}\}$ is the same as menu $\lambda\{\text{school 1}, \text{school 2}\} + (1 - \lambda)\{\text{school 2}\}$.

²⁸For expositional simplicity, Axiom 3 imposes all static REU conditions on the extended stochastic choice rule. However, it is worth noting that this is stronger than necessary: For each static REU condition except Mixture Continuity and Finiteness, imposing the condition only on the non-extended choice rule is enough to ensure (by Definition 3) that it is also satisfied by the extended choice rule.

Axiom 3 (History-dependent REU). Both ρ_0 and $\rho_1(\cdot|h^0)$ for each h^0 satisfy Axiom 0.²⁹

Our final axiom reflects the way in which tie-breaking can affect the observed choice distribution. We first define menus and histories without ties directly from choice behavior. The idea is that menus without ties are characterized by the fact that slightly perturbing their elements has no effect on choice probabilities.³⁰ We capture such perturbations using convergence in mixture, as defined following Axiom 0.

Definition 4. The set of period-0 *menus without ties*, denoted \mathcal{A}_0^* , consists of all $A_0 \in \mathcal{A}_0$ such that for any $p_0 \in A_0$ and any sequences $p_0^n \rightarrow^m p_0$ and $B_0^n \rightarrow^m A_0 \setminus \{p_0\}$, we have

$$\lim_{n \rightarrow \infty} \rho_0(p_0^n; B_0^n \cup \{p_0^n\}) = \rho_0(p_0; A_0).$$

The set of period 0 *histories without ties* is $\mathcal{H}_0^* := \{h^0 = (A_0, p_0) \in \mathcal{H}_0 : A_0 \in \mathcal{A}_0^*\}$.

The following axiom relates choice distributions after nearby histories. To state this formally, we extend convergence in mixture to histories: We say $h^{0,n} \rightarrow^m h^0$ if $h^{0,n} = (A_0^n, p_0^n)$ and $h^0 = (A_0, p_0)$ satisfy $A_0^n \rightarrow^m A_0$ and $p_0^n \rightarrow^m p_0$.

Axiom 4 (History Continuity). For all A_1, p_1 , and h^0 ,

$$\rho_1(p_1; A_1|h^0) \in \text{co}\{\lim_n \rho_1(p_1; A_1|h^{0,n}) : h^{0,n} \rightarrow^m h^0 \text{ and } h^{0,n} \in \mathcal{H}_0^*\}.$$

In general, if period-0 histories are slightly altered, we expect subsequent period-1 choice behavior to be adjusted continuously, except when there was tie-breaking in the past. If the agent chose p_0 from A_0 as a result of tie-breaking, then slightly altering the choice problem can change the set of states at which p_0 would be chosen and hence lead to a discontinuous change in the private information revealed by the choice of p_0 . The history continuity condition restricts the types of discontinuities ρ_1 can admit, ruling out situations in which choices after some history are completely unrelated to choices after any nearby history. Specifically, the fact that choice behavior after h^0 can be expressed as a mixture of behavior after some nearby histories without ties reflects the way in which the agent's tie-breaking procedures may vary with her payoff-relevant private information.

²⁹Lemma E.1 verifies that each X_t ($t = 0, 1$) is a separable metric space. Then Mixture Continuity and Finiteness make use of the same convergence notions as defined following Axiom 0.

³⁰Lu (2016b) and Lu and Saito (2018a) use an alternative approach, directly incorporating into the primitive a collection of measurable sets that capture the absence of ties and defining choice probabilities only on measurable subsets of each menu. Their approach requires that ties occur with probability either zero or one, so is not applicable to our setting. Our perturbation-based approach is similar in spirit to Ahn and Sarver (2013).

3.4 Representation Theorem

Theorem 1. Suppose that $T = 1$. Then the dynamic stochastic choice rule ρ satisfies Axioms 1–4 if and only if ρ admits a DREU representation.

The proof of Theorem 1 appears in Appendix B. We now sketch the argument for sufficiency. Readers wishing to proceed directly to the analysis of Bayesian evolving utility and evolving beliefs may skip ahead to Section 4.

First, Axiom 3 together with Theorem 0 yields a static REU representation $\mathcal{R}_0 = (\Omega_0, \mathcal{F}_0^*, \mu_0, \mathcal{F}_0, U_0, W_0)$ of ρ_0 . For each $h^0 \in \mathcal{H}_0$, Axiom 3 and Theorem 0 also imply that $\rho_1(\cdot|h^0)$ admits a static REU representation, but without ensuring any relationship between the period-0 and period-1 representations. By contrast, DREU requires that \mathcal{R}_0 be extended to a representation on a single probability space Ω, μ such that $\rho_1(p_1; A_1|A_0, p_0)$ is the conditional probability of the event $C(p_1, A_1)$ given the event $C(p_0, A_0)$.

To obtain such a representation, we only construct static REU representations of ρ_1 following specific histories that *uniquely* reveal the agent’s period-0 utility. Concretely, by simplicity of (U_0, \mathcal{F}_0) , there are finitely many possible realizations U_0^1, \dots, U_0^n of the agent’s period-0 utility, where all U_0^i are nonconstant and ordinally distinct. Thus, standard arguments (Lemma E.2) yield a menu $D_0 = \{q_0^i : i = 1, \dots, n\}$ that strictly *separates* period-0 utilities, in the sense that each q_0^i is chosen from D_0 precisely when the agent’s utility is U_0^i ; that is, the event $C_0(q_0^i, D_0)$ in Ω_0 equals the event $\{U_0 = U_0^i\}$. Figure 4 illustrates. Let $\mathcal{R}_1^i = (\Omega_1^i, \mathcal{F}_1^{*i}, \mu_1^i, \mathcal{F}_1^i, U_1^i, W_1^i)$ be a static REU representation of $\rho_1(\cdot|D_0, q_0^i)$.

The key step is to combine \mathcal{R}_0 and \mathcal{R}_1^i into a representation of ρ_1 following arbitrary histories (A_0, p_0) . Specifically, we show that for any p_1 and A_1 ,

$$\rho_1(p_1; A_1|A_0, p_0) = \sum_{i=1}^n \mu_1^i(C_1^i(p_1, A_1)) \mu_0(\{U_0 = U_0^i\}|C_0(p_0, A_0)), \quad (5)$$

where $C_1^i(p_1, A_1)$ is the event in Ω_1^i that p_1 is chosen from A_1 and $C_0(p_0, A_0)$ is the event in Ω_0 that p_0 is chosen from A_0 . Given (5), it is then straightforward to combine \mathcal{R}_0 and \mathcal{R}_1^i into a DREU representation of ρ .³¹

The argument for (5) proceeds in two steps. First, Lemma B.3 establishes (5) for histories (A_0, p_0) that are only consistent with a *single* period-0 utility U_0^i ; that is, $\mu_0(\{U_0 = U_0^i\}|C_0(p_0, A_0)) = 1$ for some i . To see the idea, note that when (A_0, p_0) is a history without ties, (A_0, p_0) and (D_0, q_0^i) reveal exactly the same information about period-0 private information. Given this, Lemma B.3 applies the two history independence conditions, Axioms 1

³¹Specifically, let $\Omega := \bigcup_{i=1}^n \{\omega_0 \in \Omega_0 : U_0(\omega_0) = U_0^i\} \times \Omega_1^i$ and define μ on Ω by $\mu(E_0 \times E_1) = \mu_0(E_0) \times \mu_1^i(E_1)$ for any events $E_0 \subseteq \{U_0 = U_0^i\}$ and $E_1 \subseteq \Omega_1^i$. If filtrations, utilities, and tie-breakers on Ω are induced from \mathcal{R}_0 and \mathcal{R}_1^i in the natural way, then (5) implies (3), as required.




Figure 4: Suppose the possible period-0 utilities are U_0^1, U_0^2, U_0^3 . Menu D_0 is a separating menu from which q_0^i is chosen precisely if $U_0 = U_0^i$. In menu $A_0 = \{p_0, r_0\}$, p_0 is chosen with probability 1 if $U_0 = U_0^1$; tied with r_0 if $U_0 = U_0^2$; and never chosen if $U_0 = U_0^3$. In $\hat{A}_0 = \frac{1}{2}A_0 + \frac{1}{2}D_0$, p_0 is replaced with three copies $p_0^i = \frac{1}{2}p_0 + \frac{1}{2}q_0^i$ with the property that $C_0(p_0^i, \hat{A}_0) = C_0(p_0, A_0) \cap \{U_0 = U_0^i\}$.

and 2, to show that $\rho_1(\cdot | D_0, q_0^i) = \rho_1(\cdot | A_0, p_0)$ coincide. Moreover, using History Continuity, the argument extends even when (A_0, p_0) features ties.

Second, Lemma B.4 establishes (5) for *arbitrary* histories (A_0, p_0) . The key idea is to consider the mixture $\hat{A}_0 := \frac{1}{2}A_0 + \frac{1}{2}D_0$ of A_0 with the separating menu D_0 . In \hat{A}_0 , p_0 is replaced with n “copies,” $p_0^i := \frac{1}{2}p_0 + \frac{1}{2}q_0^i$ for $i = 1, \dots, n$; see Figure 4. By Linear History Independence and the definition of ρ_1 following a set of histories, we have

$$\rho_1(p_1; A_1 | A_0, p_0) = \rho_1(p_1; A_1 | \hat{A}_0, \frac{1}{2}\{p_0\} + \frac{1}{2}D_0) = \sum_{i=1}^n \rho_1(p_1; A_1 | \hat{A}_0, p_0^i) \frac{\rho_0(p_0^i; \hat{A}_0)}{\sum_{j=1}^n \rho_0(p_0^j; \hat{A}_0)}. \quad (6)$$

But note that, as illustrated in Figure 4, $p_0^i = \frac{1}{2}p_0 + \frac{1}{2}q_0^i$ is chosen from $\hat{A}_0 = \frac{1}{2}A_0 + \frac{1}{2}D_0$ in precisely those states of the world where p_0 is chosen from A_0 and q_0^i is chosen from D_0 ; that is, $C_0(p_0^i, \hat{A}_0) = C_0(p_0, A_0) \cap C_0(q_0^i, D_0)$. Since by construction of the separating menu D_0 , we have $C_0(q_0^i, D_0) = \{U_0 = U_0^i\}$, this implies $\rho_0(p_0^i; \hat{A}_0) = \mu_0(C_0(p_0, A_0) \cap \{U_0 = U_0^i\})$. Moreover, when $\rho_0(p_0^i; \hat{A}_0) > 0$, then the previous paragraph (Lemma B.3) yields $\rho_1(p_1; A_1 | \hat{A}_0, p_0^i) = \mu_1^i(C_1^i(p_1, A_1))$. Combining these observations with (6) and applying Bayes’ rule yields (5), as required.

4 Characterization of BEU and BEB

DREU imposes no discipline on how the agent evaluates continuation problems. We now build on the characterization of DREU by introducing axioms that capture the dynamic sophistication of Bayesian rational agents: Section 4.1 characterizes Bayesian evolving utility (BEU), and Section 4.2 captures the additional behavioral content of its special case, Bayesian evolving beliefs (BEB). These characterizations serve as a basis for Section 5, where we contrast BEU

with dynamic discrete choice models. For simplicity of exposition, we again present our axioms in the two-period setting ($T = 1$); generalizations to an arbitrary finite horizon are provided in Appendices C–D.

4.1 Bayesian Evolving Utility

BEU is characterized by the following three axioms. First, Separability ensures that period-0 utility in every state has an additively separable form $U_0(z_0, A_1) = u_0(z_0) + V_0(A_1)$:

Axiom 5 (Separability). For any A_0 and $p_0, q_0 \notin A_0$ with $p_0^Z = q_0^Z$, $p_0^A = q_0^A$, and $A_0 \cup \{p_0\}, A_0 \cup \{q_0\} \in \mathcal{A}_0^*$, we have $\rho_0(p_0; A_0 \cup \{p_0\}) = \rho_0(q_0; A_0 \cup \{q_0\})$.

Axiom 5 is a stochastic-choice analog of the standard separability axiom for deterministic preferences (e.g., Fishburn, 1970), which requires that the agent does not care about how today’s consumption and tomorrow’s menu are correlated. That is, they do not distinguish between lotteries p_0 and q_0 that share the same marginals over both today’s consumption and tomorrow’s menu.³²

The next axiom adapts conditions from Dekel, Lipman, and Rustichini (2001) to a stochastic-choice setting, to ensure that $V_0(A_1)$ captures the option value contained in menu A_1 , i.e., that $V_0(A_1) = \mathbb{E}[\max_{p_1 \in A_1} \hat{U}_1(p_1) \mid \mathcal{F}_0]$ for some random utility function \hat{U}_1 . For part (ii), we let $m_1, m'_1 \in \Delta(\mathcal{A}_1)$ denote typical distributions over period-1 menus, and for each such m_1 , we let $\bar{A}(m_1)$ denote the average menu induced by m_1 ; that is, $\bar{A}(m_1) := \sum_{A_1 \in \mathcal{A}_1} m_1(A_1)A_1$.

Axiom 6 (Stochastic DLR).

(i). *Preference for Flexibility*: For any A_1, B_1 such that $A_1 \subseteq B_1$ and $\{(z, A_1), (z, B_1)\} \in \mathcal{A}_0^*$,

$$\rho_0((z, B_1); \{(z, A_1), (z, B_1)\}) = 1.$$

(ii). *Reduction of Mixed Menus*: For any A_0 and $(z, m_1), (z, m'_1) \notin A_0$ such that $\bar{A}(m_1) = \bar{A}(m'_1)$ and $A_0 \cup \{(z, m_1)\}, A_0 \cup \{(z, m'_1)\} \in \mathcal{A}_0^*$, we have

$$\rho_0((z, m_1); A_0 \cup \{(z, m_1)\}) = \rho_0((z, m'_1); A_0 \cup \{(z, m'_1)\}).$$

(iii). *Continuity*: $\rho_0 : \mathcal{A}_0^* \rightarrow \Delta(\Delta(X_0))$ is continuous.

(iv). *Menu Nondegeneracy*: $\{(z, A_1), (z, B_1)\} \in \mathcal{A}_0^*$ for some z, A_1, B_1 .

³²Lu and Saito (2018b) study a random utility model where separability is violated, as in Epstein and Zin (1989). They show that even on simple domains where the continuation menu is fixed the analyst’s estimates of the function u are biased because they are contaminated by the nonlinear continuation utility.

Part (i) corresponds to [Kreps’s \(1979\)](#) “preference for flexibility” axiom, which says that the agent always (weakly) prefers bigger menus. This captures a key property of Bayesian-rational agents in a dynamic setting, viz., a positive option value. The axiom is violated in [Ke’s \(2018\)](#) model, where the agent has a deterministic utility but anticipates making random execution mistakes. This agent’s choices over menus exhibit a form of preference for commitment, because eliminating inferior options from a menu benefits the agent by reducing the scope for mistakes. Preference for flexibility is also violated by dynamic logit ([Fudenberg and Strzalecki, 2015](#)) and more general dynamic discrete choice models, as we will discuss in more detail in Section 5. Part (ii) requires that the agent reduces mixtures over menus; this is analogous to Menu Independence in [Dekel, Lipman, and Rustichini \(2001\)](#) and implies that the agent cannot affect tomorrow’s utility distribution. Parts (iii) and (iv) ensure that the agent has continuous and nontrivial preferences over continuation menus.

The final axiom adapts the sophistication axiom due to [Ahn and Sarver \(2013\)](#). Fix any history $h^0 = (A_0, p_0)$ and menus $B_1 \supset A_1$. We require that if the agent sometimes chooses an option in $B_1 \setminus A_1$ following history h^0 , then in some states of the world in which she chooses p_0 from A_0 , she must value menu B_1 strictly more than A_1 (and vice versa). This axiom ensures that the agent correctly anticipates her future utility distribution; that is, $\hat{U}_1 = U_1$.

To formalize this, we must express in terms of stochastic choices the fact that in some states of the world in which the agent chooses p_0 from A_0 , she values menu B_1 strictly more than A_1 . This goes beyond [Ahn and Sarver \(2013\)](#), whose setting in period 0 features no consumption and no randomness in the agent’s preference over period-1 menus.³³ To see the idea, suppose that for some lotteries q_0 and r_0 , we have

$$\rho_0 \left(\frac{1}{2}p_0 + \frac{1}{2}q_0; \frac{1}{2}A_0 + \frac{1}{2}\{q_0, r_0\} \right) > 0. \tag{7}$$

Then we can conclude that in some states of the world in which the agent chooses p_0 from A_0 , she *weakly* prefers q_0 to r_0 : Indeed, for an expected-utility maximizer, it is optimal to choose $\frac{1}{2}p_0 + \frac{1}{2}r_0$ from menu $\frac{1}{2}A_0 + \frac{1}{2}\{q_0, r_0\}$ if and only if it is optimal to choose p_0 from A_0 *and* to choose r_0 from $\{q_0, r_0\}$.³⁴ To be able to infer that in some states where the agent chooses p_0 from A_0 she *strictly* prefers q_0 to r_0 , we must additionally ensure that in some such states the menu $\{q_0, r_0\}$ does not feature a tie. Similar to [Ahn and Sarver \(2013\)](#), this is achieved by

³³Likewise, no such challenge arises in [Fudenberg and Strzalecki’s \(2015\)](#) dynamic logit model. Because of their i.i.d. shocks assumption, the agent’s preference over continuation menus does not vary with her period-0 consumption choices.

³⁴This observation is related to the random incentive mechanism used in experimental work. To elicit a subject’s ranking over a number of options in an incentive compatible manner, the subject is asked to indicate choices from multiple menus; a lottery then determines which menu (and corresponding choice) is implemented. See e.g., [Becker, DeGroot, and Marschak \(1964\)](#) and [Chambers and Lambert \(2017\)](#).

requiring (7) to hold for all small enough perturbations $q_0^n \rightarrow^m q_0$ and $r_0^n \rightarrow^m r_0$.³⁵

Point (ii) of the following axiom applies this idea with $q_0 = (z, B_1)$ and $r_0 = (z, A_1)$ for an arbitrary consumption z .

Axiom 7 (Sophistication). For any $h^0 = (A_0, p_0) \in \mathcal{H}_0^*$, z , and $A_1 \subseteq B_1 \in \mathcal{A}_1^*(h^0)$, the following are equivalent:

- (i). $\rho_1(p_1; B_1 | h^0) > 0$ for some $p_1 \in B_1 \setminus A_1$
- (ii). $\liminf_n \rho_0(\frac{1}{2}p_0 + \frac{1}{2}(z, B_1^n); \frac{1}{2}A_0 + \frac{1}{2}\{(z, B_1^n), (z, A_1^n)\}) > 0$ for all $A_1^n \rightarrow^m A_1, B_1^n \rightarrow^m B_1$.

Axiom 7 applies only to menus B_1 that do not feature ties conditional on history h^0 . Analogous to Definition 4, for any $h^0 \in \mathcal{H}_0$, the set of period-1 *menus without ties conditional on history* h^0 is denoted $\mathcal{A}_1^*(h^0)$ ³⁶ and consists of all $A_1 \in \mathcal{A}_1$ such that for any $p_1 \in A_1$ and any sequences $p_1^n \rightarrow^m p_1$ and $B_1^n \rightarrow^m A_1 \setminus \{p_1\}$, we have $\lim_{n \rightarrow \infty} \rho_1(p_1^n; B_1^n \cup \{p_1^n\} | h^0) = \rho_1(p_1; A_1 | h^0)$.

Theorem 2. Suppose $T = 1$ and ρ admits a DREU representation. Then ρ satisfies Axioms 5–7 if and only if ρ admits a BEU representation.

Proof. See Appendix C. ■

4.2 Bayesian Evolving Beliefs

Bayesian evolving beliefs is a specialization of BEU where the agent has a time-invariant but unknown felicity about which she learns over time. In this section, we characterize the additional behavioral content of this assumption by a simple axiom on the agent’s choices over streams of consumption lotteries. Section 6.2 provides an alternative characterization on the subdomain where in each period the agent chooses only today’s consumption.

Given consumption lotteries $\ell_0, \ell_1 \in \Delta(Z)$, let the *stream* $(\ell_0, \ell_1) \in \Delta(X_0)$ be the period-0 lottery that in period 1 yields consumption lottery ℓ_1 for sure and in period 0 yields consumption according to ℓ_0 ; formally, $(\ell_0, \ell_1) = p_0$ where $p_0^Z = \ell_0$ and $p_0^A = \delta_{\{\ell_1\}}$.

Axiom 8 (Stationary Consumption Preference). If $(\ell, \ell), (\ell', \ell') \in A_0 \in \mathcal{A}_0^*$, then $\rho_0((\ell, \ell'); A_0) = 0$.

³⁵In an earlier working paper version, we apply this idea more generally to define an incomplete and history-dependent revealed preference relation \succsim_{h^t} that captures that one lottery is preferred to another in any state of the world ω that gives rise to history h^t ; see Section 4.1 of Frick, Iijima, and Strzalecki (2017). This preference relation can be used to provide alternative versions of Axioms 5–8.

³⁶Note that $\mathcal{A}_1^*(h^0) \not\subseteq \mathcal{A}_1(h^0)$ because the first set contains all menus without ties (we use history h^0 here only to determine where ties could occur) while the second set contains only menus that occur with positive probability after history h^0 —typically very few menus.

Axiom 8 requires that the agent never chooses to commit to a time-varying consumption stream (ℓ, ℓ') if her choice set also contains the corresponding stationary consumption streams (ℓ, ℓ) and (ℓ', ℓ') . This reflects Bayesian learning about fixed but unknown tastes: Indeed, if the agent currently believes ℓ to be better than ℓ' , then by the martingale property of beliefs she should expect her information next period to still favor ℓ on average and will hence prefer (ℓ, ℓ) to (ℓ, ℓ') (and analogously in the opposite case).

To characterize BEB, we postulate the existence of a pair $\bar{\ell}, \underline{\ell}$ of consumption lotteries such that the agent always strictly prefers $\bar{\ell}$ to $\underline{\ell}$ at all times and histories. This condition is innocuous if, for example, the outcome space includes a monetary dimension.

Condition 1 (Uniformly Ranked Pair). There exist $\bar{\ell}, \underline{\ell} \in \Delta(Z)$ such that for all $\ell \in \Delta(Z)$ and h^0 , we have $A_0 := \{(\bar{\ell}, \ell), (\underline{\ell}, \ell)\} \in \mathcal{A}_0^*$, $A_1 := \{\bar{\ell}, \underline{\ell}\} \in \mathcal{A}_1^*(h^0)$, and $\rho_0((\bar{\ell}, \ell); A_0) = \rho_1(\bar{\ell}; A_1 | h^0) = 1$.

Theorem 3. Suppose that $T = 1$ and ρ admits a BEU representation and satisfies Condition 1. Then ρ satisfies Axiom 8 if and only if ρ admits a BEB representation.

The proof is in Appendix D. The key idea is to show that Axiom 8 is equivalent to the requirement that $u_0(\omega)$ and $\mathbb{E}[u_1 | \mathcal{F}_0(\omega)]$ represent the same preference over consumption lotteries in all states, which after appropriate normalization yields (2).

We note that while BEB allows for a stochastic process $\delta_t : \Omega \rightarrow \mathbb{R}_{++}$ of discount factors, an earlier working paper version of this article includes an additional axiom that ensures a deterministic discount factor $\delta > 0$; moreover, a standard impatience axiom corresponds to $\delta < 1$.³⁷

Remark 1 (Identification). Proposition I.1 in Appendix I.1 establishes identification results for DREU, BEU, and BEB. To summarize, the identification result for DREU is a period-by-period analog of the known result for static REU (Proposition 4 in Ahn and Sarver (2013)); that is, ρ uniquely determines the underlying stochastic process of *ordinal* payoff-relevant private information and the (ordinal) distribution of tie-breakers for choices featuring ties. The result for BEU generalizes Theorem 2 of Ahn and Sarver (2013), implying strictly sharper identification than DREU of the agent's *cardinal* private information. In particular, BEU allows for meaningful intertemporal comparisons of utility in each state and for limited cross-state comparisons of utility within states that correspond to the same period-0 private information. Finally, BEB, unlike BEU, allows for unique identification of the discount factor process and entails even sharper identification of cardinal private information.³⁸ ▲

³⁷See Frick, Iijima, and Strzalecki (2017) Axiom 9 for the former and p. 26 for the latter.

³⁸The discount factor process is unique in other special cases of BEU as well; for example if each alternative z consists of wealth and a consumption bundle and the utility of wealth is separable and state-independent.

5 Comparison with Dynamic Discrete Choice

In this section, we compare Bayesian evolving utility to dynamic discrete choice (DDC) models that are widely used in empirical work. The key distinction we highlight concerns the way in which random utility shocks are modeled: BEU is a special case of the most general DDC model, but while BEU features only *shocks to consumption*, most DDC models introduce more general *shocks to actions*. We show that under certain widely used assumptions, the latter form of shocks leads to violations of a key feature of Bayesian rationality, namely positive option value. We also illustrate how this can lead to biased parameter estimates.

5.1 DDC Models

For simplicity, we restrict the domain to *deterministic* decision trees, where each period- t outcome space Y_t consists of pairs $y_t = (z_t, A_{t+1})$ of instantaneous consumptions z_t and continuation menus A_{t+1} . We refer to each y_t as an *action*.

The following special case of DREU encompasses many models in the dynamic discrete choice literature (for surveys, see [Aguirregabiria and Mira, 2010](#); [Rust, 1994](#)):³⁹

Definition 5. The *DDC model* is a restriction of DREU to deterministic decision trees that additionally satisfies the Bellman equation

$$U_t(z_t, A_{t+1}) = v_t(z_t) + \delta \mathbb{E} \left[\max_{y_{t+1} \in A_{t+1}} U_{t+1}(y_{t+1}) | \mathcal{F}_t \right] + \varepsilon_t^{(z_t, A_{t+1})}, \quad (8)$$

with deterministic felicities $v_t : Z \rightarrow \mathbb{R}$, \mathcal{F}_t -adapted zero-mean *shocks to actions* $\varepsilon_t : \Omega \rightarrow \mathbb{R}^{Y_t}$, and a discount factor $\delta \in (0, 1)$.⁴⁰

Observe that BEU corresponds precisely to the special case of DDC where the ε shocks do not apply to general actions $y_t = (z_t, A_{t+1})$, but only to instantaneous consumptions z_t ; formally, in all periods t , any actions (z_t, A_{t+1}) and (z_t, B_{t+1}) that feature the same consumption z_t receive

³⁹For ease of comparison with BEU, we impose the following three restrictions, which are extraneous to the distinction between shocks to consumption and shocks to actions that we highlight in this section. First, we impose [Rust's \(1994\)](#) assumption AS, viz. that ε shocks enter into (8) in an additively separable manner; this is widely, but not universally imposed in the DDC literature. Violations of AS can be accommodated by DREU, but are incompatible with BEU in ways that are orthogonal to the focus of this section. Second, whereas filtration \mathcal{F}_t is exogenous, DDC models often allow the agent's choices to affect transitions from the current state to tomorrow's state; this can be accommodated by a consumption-dependent extension of DREU (see [Section 7.2](#)). Finally, ε can capture any state variables that are privately observed by the agent, but in contrast with many DDC models (e.g., [Hotz and Miller, 1993](#)), all representations in this paper abstract away from state variables that are *jointly* observed by the analyst and agent (save for menus); [Duraj \(2018\)](#) extends DREU to incorporate the latter.

⁴⁰We assume a deterministic $\delta \in (0, 1)$ for simplicity. As under BEU, δ is not identified under general DDC, but this poses no problems in the specific examples we consider. See also footnote 52.

the same shock

$$\varepsilon_t^{(z_t, A_{t+1})} = \varepsilon_t^{(z_t, B_{t+1})} =: \varepsilon_t^{z_t}. \quad (9)$$

Indeed, given (9), setting $u_t(z_t) = v_t(z_t) + \varepsilon_t^{z_t}$ yields an \mathcal{F}_t -adapted process of felicities that satisfies (1); and conversely, given any \mathcal{F}_t -adapted felicity process u_t satisfying (1), we can let $v_t(z_t) := \mathbb{E}[u_t(z_t)]$ and $\varepsilon_t^{z_t} := u_t(z_t) - v_t(z_t)$.

Thus, Theorem 2 provides an axiomatic foundation for this *shocks to consumption* version of DDC, while Proposition I.1 is an identification result.⁴¹ Shocks to consumption alone are sufficient to capture phenomena such as permanent unobserved heterogeneity and serially correlated unobserved state variables that are studied in the DDC literature; indeed, Pakes (1986) can be viewed as an early special case of BEU. Under shocks to consumption, all randomness in the agent's evaluation of continuation menus A_{t+1} is captured by the \mathcal{F}_t -adapted continuation value $\mathbb{E}[\max_{y_{t+1} \in A_{t+1}} U_{t+1}(y_{t+1}) | \mathcal{F}_t]$ that reflects the agent's private information about future shocks to consumption.

However, for estimation purposes, many central models in the DDC literature introduce shocks to actions that violate (9) by applying additional shocks to continuation menus that may be completely detached from their continuation value. The main purpose of introducing such general shocks to actions is that under suitable assumptions they can ensure *nondegenerate likelihoods*: This denotes the property that in any menu A_t , all actions $y_t \in A_t$ are chosen with positive probability at all histories, which is central for statistical estimation. One of the most widely used models with this property is the following i.i.d. version of DDC:⁴²

Definition 6. The *i.i.d. DDC model* is a restriction of DDC such that

$$U_t(z_t, A_{t+1}) = v_t(z_t) + \delta \mathbb{E} \left[\max_{y_{t+1} \in A_{t+1}} U_{t+1}(y_{t+1}) \right] + \varepsilon_t^{(z_t, A_{t+1})},$$

where for all periods t and τ and all actions (z_t, A_{t+1}) and $(x_\tau, B_{\tau+1})$, $\varepsilon_t^{(z_t, A_{t+1})}$ and $\varepsilon_\tau^{(x_\tau, B_{\tau+1})}$ are independently and identically distributed random variables with a full support density.⁴³

Under i.i.d. DDC, both felicities v_t and continuation values $\mathbb{E}[\max_{y_{t+1} \in A_{t+1}} U_{t+1}(y_{t+1})]$ are deterministic, and all randomness in the agent's evaluation of z_t and A_{t+1} is fully captured by $\varepsilon_t^{(z_t, A_{t+1})}$. Since these shocks are i.i.d. across any pairs of *actions*, including actions (z_t, A_{t+1}) and (z_t, B_{t+1}) that differ only in their continuation menus, they violate (9). The following

⁴¹Our identification result is complementary to those in econometrics (Hu and Shum, 2012; Kasahara and Shimotsu, 2009; Magnac and Thesmar, 2002; Norets and Tang, 2013), because we allow for menu variation but abstract from jointly observable state variables.

⁴²See e.g., Miller (1984), Rust (1989), Hendel and Nevo (2006), Kennan and Walker (2011), Sweeting (2013), and Gowrisankaran and Rysman (2012).

⁴³While DREU assumes finitely generated distributions, a full support density distribution is observationally equivalent to one with a sufficiently large finite support given the finiteness of the deterministic decision tree domain.

section shows that this leads to behavior that is incompatible with any BEU model, including the i.i.d. version of BEU where ε shocks satisfy (9) and $\varepsilon_t^{z_t}$ and $\varepsilon_\tau^{x_\tau}$ are i.i.d. across all pairs of *consumptions*.

5.2 Option Value in BEU vs. i.i.d. DDC

We now show that, in contrast with BEU, shocks to actions that violate (9) can lead to behavior that displays a negative option value. To make this point most clearly, we focus predominantly on i.i.d. DDC, before briefly turning to more general models. Given that i.i.d. DDC is a workhorse model for structural estimation, understanding its properties is also important in its own right.

The first manifestation of a negative option value is that the i.i.d. DDC agent sometimes chooses to commit to strictly smaller menus. Suppose there are two periods, $t = 0, 1$. Let $A_0 := \{(z_0, A_1^{\text{small}}), (z_0, A_1^{\text{big}})\}$ where $A_1^{\text{small}} = \{z_1\}$ and $A_1^{\text{big}} = \{z_1, z'_1\}$. From Axiom 6 it follows that under BEU, $\rho_0((z_0, A_1^{\text{small}}), A_0) = 0$ absent ties. By contrast, under i.i.d. DDC this probability is strictly positive.

Proposition 1. Under i.i.d. DDC, we have $0 < \rho_0((z_0, A_1^{\text{small}}); A_0) < 0.5$. Moreover, if the ε shocks are scaled by $\lambda > 0$, then $\rho_0((z_0, A_1^{\text{small}}); A_0)$ is strictly increasing in λ whenever $v_1(z'_1) > v_1(z_1)$.

All proofs for this section appear in Supplementary Appendix G. The first part follows from the fact that by design, i.i.d. DDC features nondegenerate likelihoods. Specifically, the agent chooses $(z_0, A_1^{\text{small}})$ from A_0 whenever the realization of $\varepsilon_0^{(z_0, A_1^{\text{small}})}$ exceeds $\varepsilon_0^{(z_0, A_1^{\text{big}})}$ by more than the expected utility difference of the two menus, and since the two shocks are i.i.d. with full support, this happens with strictly positive probability. Nevertheless, since $\mathbb{E}[U_0(z_0, A_1^{\text{big}})] > \mathbb{E}[U_0(z_0, A_1^{\text{small}})]$, this probability is less than 0.5. The second part of Proposition 1 further highlights the negative effect of i.i.d. shocks to actions on option value by showing that greater variance in ε can increase the probability of choosing the small menu, even though this increases the continuation value of the larger menu.⁴⁴

More strikingly, there are decision problems for which behavior under i.i.d. DDC displays a negative option value with probability greater than 0.5. Specifically, consider the following decision timing problem, illustrated in Figure 5. There are three periods $t = 0, 1, 2$. The consumption in period 2 is either y or z , depending on the agent's choice. The agent can make her decision early, committing in period 1 to receiving y or z the following period; or she can make the decision late, maintaining full flexibility about choosing y or z until the final period. The decision when to choose is made in period 0, and the consumption in periods 0 and 1

⁴⁴We thank Jay Lu for suggesting that we investigate this comparative static.




Figure 5: Decision timing.

is x irrespective of the agent's decision; for simplicity assume that the utility of x is always zero. To fix ideas, assume that a student was admitted to two PhD programs (y and z) and is considering whether to make her decision before the visit days ($t = 1$) or after ($t = 2$); assume that she plans to attend the visit days regardless. Formally, in period 0 the agent faces the menu $A_0 = \{(x, A_1^{\text{early}}), (x, A_1^{\text{late}})\}$, and in period 1 she faces either menu $A_1^{\text{early}} = \{(x, \{y\}), (x, \{z\})\}$ or menu $A_1^{\text{late}} = \{(x, \{y, z\})\}$, depending on her period-0 choice.

Proposition 2. Under BEU, $\rho_0((x, A_1^{\text{early}}); A_0) = 0$ absent ties. Under i.i.d. DDC, $\rho_0((x, A_1^{\text{early}}); A_0) > 0.5$; moreover, if ε is scaled by $\lambda > 0$, then $\rho_0((x, A_1^{\text{early}}); A_0)$ is strictly increasing in λ whenever $v_2(y) \neq v_2(z)$.

In this decision problem there is no penalty to deciding late, as the timing of the decision does not affect the timing of the consumptions y or z . Thus, reflecting a positive option value, the BEU agent chooses to make decisions late because waiting until the final period enables her to better tailor her choice to her realized felicity. This prediction does not rely on serially correlated private information; indeed, it remains true under i.i.d. BEU.

By contrast, Proposition 2 shows that the i.i.d. DDC agent chooses to decide *early* with probability greater than 0.5. To see why, consider the simplest case when $v_2(y) = v_2(z)$. In this case, the choice boils down to comparing $\delta \mathbb{E}[\max\{\varepsilon_1^{(x, \{y\})}, \varepsilon_1^{(x, \{z\})}\}]$ and $\delta^2 \mathbb{E}[\max\{\varepsilon_2^y, \varepsilon_2^z\}]$. Since the ε shocks are i.i.d. and mean zero and $\delta \in (0, 1)$, the former dominates the latter, so that the agent chooses to decide early with probability greater than 0.5. Intuitively, choosing early is attractive because it allows the agent to obtain a positive payoff, namely the maximum of two i.i.d. mean zero shocks, early while deferring the choice delays those payoffs. Again, the negative effect of the ε shocks on option value is further reflected by the fact that the agent's preference for deciding early is increasing in their variance, even though this increases uncertainty about future payoffs.

A special case of the preference for early decisions under i.i.d. logit ε shocks was proved

by Fudenberg and Strzalecki (2015), by examining the closed-form expressions for continuation values in this setting.⁴⁵ Proposition 2 shows that this result does not rely on those specific expressions. Rather, it is a consequence of the mechanical nature of shocks to actions in *any* i.i.d. DDC model: As we discussed above, unlike shocks to consumption, these shocks apply directly to continuation menus in a way that is completely detached from their expected continuation value.⁴⁶

Finally, we note that the findings in this section are not limited to i.i.d. DDC. Indeed, the following two widely studied DDC models depart, respectively, from the assumption that shocks are i.i.d. over time or i.i.d. across actions, but continue to display a preference for early decisions.⁴⁷ First, under DDC with *permanent unobserved heterogeneity*, ε displays the following form of serial correlation: Each shock $\varepsilon_t^{(z_t, A_{t+1})} = \pi_t^{z_t} + \theta_t^{(z_t, A_{t+1})}$ is decomposed into a “permanent” shock $\pi_t^{z_t}$ that is measurable with respect to \mathcal{F}_0 and a “transitory” shock $\theta_t^{(z_t, A_{t+1})}$ that conditional on \mathcal{F}_0 is i.i.d. across all periods and actions. Thus, utility in each period depends on the agent’s “type” (which she learns in period 0), but each type of agent is also subject to i.i.d. shocks to actions. In this model, behavior ρ is a mixture of i.i.d. DDC choice rules. Thus, since Proposition 2 applies to each of these choice rules, their mixture ρ continues to satisfy $\rho_0((x, A_1^{\text{early}}); A_0) > 0.5$. Second, some models feature *transitory but correlated* shocks to actions: Here $\varepsilon_t^{(z_t, A_{t+1})}$ and $\varepsilon_\tau^{(x_\tau, B_{\tau+1})}$ are i.i.d. whenever $t \neq \tau$, but might be correlated within any fixed period $t = \tau$; e.g., due to transitory health shocks that affect the agent’s evaluation of all actions in a given period. As long as within-period shocks are not perfectly correlated, we again have $\rho_0((x, A_1^{\text{early}}), A_0) > 0.5$; intuitively, $\mathbb{E}[\max\{\varepsilon_1^{(x, \{y\})}, \varepsilon_1^{(x, \{z\})}\}] = \mathbb{E}[\max\{\varepsilon_2^y, \varepsilon_2^z\}]$ remains strictly positive, so the agent again prefers to receive this shock early.⁴⁸

5.3 Parameter Estimates in a Stopping Problem

Unlike the previous decision timing problem, many economic decisions, such as stopping problems, feature a tradeoff between an immediate payoff today and the option value of delay. We

⁴⁵Fudenberg and Strzalecki (2015) also introduced a choice-aversion parameter that scales the desire for flexibility and for early decisions. However, in this model the parameter values that imply choice of late decisions with probability greater than 0.5 also imply choice of smaller menus with probability greater than 0.5, thus making violations of positive option value particularly stark in the latter dimension.

⁴⁶Our critique of the mechanical nature of shocks to actions is complementary to Apesteguia and Ballester’s (2018) critique of i.i.d. ε in *static* settings, but the logic of our results is quite different, both formally and conceptually. In particular, in Propositions 1 and 2 these shocks lead to counterintuitive predictions at an *absolute* level, rather than only at a comparative level as in their results. Moreover, our comparative results are also quite different, as we vary the variance of ε , whereas they vary the curvature of the utility function.

⁴⁷Both permanent unobserved heterogeneity and transitory shocks that are correlated across actions are central ingredients of what Aguirregabiria and Mira (2010) (p. 42 ff.) term Eckstein-Keane-Wolpin models.

⁴⁸Predictions under even more general models are ambiguous. E.g., suppose $\varepsilon_t^{(z_t, A_{t+1})} = \pi_t^{z_t} + \theta_t^{(z_t, A_{t+1})}$ is decomposed into \mathcal{F}_t -adapted shocks to consumption $\pi_t^{z_t}$ and i.i.d. shocks to actions $\theta_t^{(z_t, A_{t+1})}$, but $\pi_t^{z_t}$ need not be \mathcal{F}_0 -measurable. This yields a hybrid of BEU and i.i.d. DDC, which may display a preference for early or late decisions depending on the relative magnitudes of π_t and θ_t and the amount of serial correlation in π_t .

now illustrate how in such settings DDC models with additional mechanical shocks to continuation menus lead to systematically different parameter estimates relative to the pure shocks to consumption model of BEU. In particular, we highlight the qualitative biases that arise if the analyst uses the former type of DDC model but the true model is BEU.⁴⁹

Consider again Example 3 from Section 1.3. In period 0, the agent chooses between two actions, consuming a today (and nothing tomorrow) or delaying consumption until period 1 where she will face menu $A_1 := \{a, b\}$. Slightly abusing notation, we denote these two period-0 actions by a and A_1 and let $A_0 := \{a, A_1\}$.⁵⁰ Let D be the set of all possible choice sequences (consume a in period 0; delay and consume a in period 1; delay and consume b in period 1). In the following, we think of the agent’s stochastic choice rule ρ as a *data generating process* over D ; that is, the analyst observes strings of data $d = (d_1, \dots, d_n) \in D^n$, where each d_i results from an independent draw according to ρ .

For concreteness, we compare parameter estimates under the following versions of i.i.d. DDC and BEU. Under i.i.d. DDC, let $v_0(a) = v_1(a) = w_a$ and $v_1(b) = w_b$ with discount factor δ . Thus, $U_1^{\text{DDC}}(x) = w_x + \varepsilon_1^x$ for $x = a, b$, $U_0^{\text{DDC}}(a) = w_a + \varepsilon_0^a$ and $U_0^{\text{DDC}}(A_1) = \delta \mathbb{E}_0[\max\{U_1(a), U_1(b)\}] + \varepsilon_0^{A_1}$, where all ε shocks are i.i.d. according to some full support distribution F with mean zero. For BEU, we consider a minimal departure from i.i.d. DDC that features the same i.i.d. shocks to instantaneous consumptions a and b : $U_1^{\text{BEU}}(x) = U_1^{\text{DDC}}(x)$ for $x = a, b$ and $U_0^{\text{BEU}}(a) = U_0^{\text{DDC}}(a)$. The only difference is that $U_0^{\text{BEU}}(A_1) = \delta \mathbb{E}_0[\max\{U_1^{\text{BEU}}(a), U_1^{\text{BEU}}(b)\}]$; i.e., there is no ε shock to the period-0 action “delay,” reflecting that this involves no instantaneous consumption.⁵¹ Let ρ^{DDC} and ρ^{BEU} denote the induced stochastic choice rules in this stopping problem.

The analyst seeks to estimate the discount factor δ and average utility difference $w := w_a - w_b$ between the two consumptions given any distribution F of i.i.d. ε shocks.⁵² To simplify notation, we normalize $w_b = 0$. Let $\Theta \subseteq \mathbb{R}^2$ denote the compact space of parameters (w, δ) that is considered by the analyst. We assume that Θ is large enough so that the data ρ is *compatible* with both models, i.e., for each $M \in \{\text{DDC}, \text{BEU}\}$, there exists $(w^M, \delta^M) \in \Theta$ such that $\rho^M = \rho$ holds under parameters (w^M, δ^M) . Let $(\hat{w}_n^M, \hat{\delta}_n^M) \in \Theta$ denote the corresponding maximum likelihood estimates under observation size n .

The following proposition shows that i.i.d. DDC tends to “exaggerate” the estimate of the discount factor relative to BEU. The result assumes that distribution F is symmetric with a

⁴⁹The quantitative importance of such biases is an empirical question, which is beyond the scope of this paper.

⁵⁰To be more precise, period-0 actions a and A_1 should be written as $(a, \{z_\emptyset\})$ and (z_\emptyset, A_1) respectively, where z_\emptyset denotes a dummy variable that corresponds to “no consumption.”

⁵¹There is another BEU specification that is observationally equivalent to i.i.d. DDC in this particular stopping problem; specifically, this version applies a shock $\varepsilon_0^{z_\emptyset}$ to the period-0 dummy consumption z_\emptyset (see footnote 50) despite the fact that z_\emptyset is only a notational stand-in for the decision to delay consumption. However, since *any* specification of BEU is incompatible with i.i.d. DDC in *some* decision trees (see Section 5.2), this model again yields different parameter estimates from i.i.d. DDC in settings other than the present stopping problem.

⁵²Unlike in general, here δ is identified for both DDC and BEU, as the average felicity of a is assumed constant across periods; this approach is also used in other stopping problems, cf. [Martinez, Meier, and Sprenger \(2017\)](#).

unimodal density (e.g., probit); Appendix G.3 shows how it generalizes to a broader class of distributions.

Proposition 3. Suppose that the data generating process ρ is compatible with both models. If F has a symmetric and unimodal density, then almost surely

(i). $\lim_n \hat{w}_n^{\text{DDC}} = \lim_n \hat{w}_n^{\text{BEU}}$

(ii). $\lim_n \hat{\delta}_n^{\text{DDC}} < \lim_n \hat{\delta}_n^{\text{BEU}}$ if $\rho_0(a; A_0) > 0.5$ and $\lim_n \hat{\delta}_n^{\text{DDC}} > \lim_n \hat{\delta}_n^{\text{BEU}}$ if $\rho_0(a; A_0) < 0.5$.

Both models yield the same estimate of w because they predict the same period-1 choice probabilities. To understand the result for δ , suppose first that $\rho_0(a, A_0) > 0.5$, i.e., the agent is more likely to choose immediate consumption than delay. Intuitively this occurs when the agent is impatient, and in this case DDC underestimates δ relative to BEU. Conversely, when the agent is patient (i.e., $\rho_0(a, A_0) < 0.5$), DDC overestimates δ relative to BEU. Thus, DDC always exaggerates the estimate of δ . The reason is precisely that DDC includes an additional mechanical shock $\varepsilon_0^{A_1}$ to the action of delaying. This creates more choice variance around modal choices in period 0; to compensate, the model must exaggerate the value difference between choices in period 0, thereby producing more extreme estimates of the discount factor.

An immediate corollary of Proposition 3 is that if the true data is in fact generated by BEU with parameters (w, δ) but the analyst uses i.i.d. DDC, then the resulting estimates almost surely satisfy (i) $\lim_n \hat{w}_n^{\text{DDC}} = w$ and (ii) $\lim_n \hat{\delta}_n^{\text{DDC}} > \delta$ if $\rho_0(a; A_0) > 0.5$ and $\lim_n \hat{\delta}_n^{\text{DDC}} < \delta$ if $\rho_0(a; A_0) < 0.5$. Finally, we note that the same logic as above can be applied to characterize the difference in estimates in other classic stopping problems, such as task completion or patent renewal.

5.4 Discussion

Our findings highlight the following modeling tradeoff. On the one hand, general shocks to actions are statistically convenient, ensuring nondegenerate likelihoods under formulations such as i.i.d. DDC, whereas BEU agents necessarily choose some options with probability 0. On the other hand, Section 5.2 shows that this convenience comes at a cost, namely significant violations of positive option value, both at an absolute and comparative level. Such violations cast doubt on the typical interpretation of ε as “unobserved utility shocks” and seem particularly problematic in applications where the modeled agents are profit-maximizing firms.⁵³

While this may seem to imply having to choose between statistical nondegeneracy and Bayesian rationality, we note that in many specific decision problems, e.g., the stopping problem

⁵³Another interpretation of ε in the DDC literature is that they capture “mistakes” or some small deviations from perfect rationality. However, Proposition 2 shows that the implied deviations are not small as they occur with probability greater than a half; moreover, this interpretation is at odds with the fact that in (8) the ε shocks enter into the agent’s expected continuation value.

in Section 5.3, versions of BEU feature nondegenerate likelihoods and can be used for parameter inference. This is also true in more concrete applications, such as in Pakes’s (1986) study of patent renewal where a BEU model is estimated. Thus, in such settings, the analyst can refrain from imposing shocks to actions and can estimate a BEU model that respects Bayesian rationality. In settings where BEU is statistically degenerate, any statistically nondegenerate model will sometimes violate Bayesian rationality, but suitable hybrid models of BEU and i.i.d. DDC (see footnote 48) may help limit the severity of the violations.

6 Atemporal Choice Domain and Choice Persistence

In this section, we restrict to the simple subdomain of *atemporal consumption problems*, where the agent chooses only (lotteries over) today’s consumption in each period and her current choices do not affect tomorrow’s menu. As illustrated in Example 1, stochastic choice data on this domain is often studied in empirical work, notably the large literature on brand choice dynamics in marketing and economics.⁵⁴ An important empirical regularity is that choice data tends to display some “persistence.” Sections 6.1 and 6.2 axiomatically characterize two notions of choice persistence, showing that they correspond precisely to two important special cases of BEU: taste persistence and learning.⁵⁵

Focusing on two periods for simplicity, our atemporal domain is formalized as follows. Given any consumption lottery $\ell_0 \in \Delta(Z)$ and menu of consumption lotteries $L_1 \in \mathcal{A}_1 = \mathcal{K}(\Delta(Z))$,⁵⁶ let (ℓ_0, L_1) denote the lottery p_0 that in period 0 yields consumption according to ℓ_0 and in period 1 yields menu L_1 for sure; that is, $p_0^Z = \ell_0$ and $p_0^A = \delta_{L_1}$. Likewise, for any menu $L_0 \in \mathcal{K}(\Delta(Z))$ of consumption lotteries and $L_1 \in \mathcal{A}_1$, define $(L_0, L_1) := \{(\ell_0, L_1) : \ell_0 \in L_0\} \in \mathcal{A}_0$ to be the menu consisting of all lotteries that yield period-0 consumption according to some $\ell_0 \in L_0$ and in period 1 yield menu L_1 for sure. Let $\mathcal{L}_0^* \subseteq \mathcal{K}(\Delta(Z))$ denote the set of consumption menus without ties, which consists of all L_0 such that $(L_0, L_1) \in \mathcal{A}_0^*$ for all $L_1 \in \mathcal{A}_1$.

We assume throughout that ρ admits a BEU representation. On our atemporal domain, this has especially simple testable implications: ρ must satisfy the restrictions to this domain of the DREU axioms (Axioms 1–4) and of Separability (Axiom 5).⁵⁷

⁵⁴E.g., Dubé, Hitsch, and Rossi (2010); Jeuland (1979); Keane (1997); Seetharaman (2004) and references therein.

⁵⁵Our characterization of the implications of choice persistence for the general BEU model is complementary to the empirical brand choice literature, which tests to what extent particular parametric or semi-parametric forms of serially correlated felicities can capture choice persistence in specific data sets. One goal of this literature is to disentangle (what we term) history dependence (e.g., persistent taste heterogeneity) and consumption dependence (e.g., habit formation) as sources of choice persistence. While the model in this section rules out consumption dependence, Section 7.2 briefly discusses how to incorporate it.

⁵⁶Throughout this section, we denote menus by L_t to emphasize that they consist of consumption lotteries.

⁵⁷Note that Axioms 6 and 7 have no bite on the atemporal domain.

Given this, we can define the restriction ρ^Z of ρ to *atemporal consumption problems without ties*: For any $L_0 \in \mathcal{L}_0^*$ and $\ell_0 \in L_0$, define $\rho_0^Z(\ell_0; L_0) := \rho_0((\ell_0, L_1); (L_0, L_1))$ for an arbitrary choice of L_1 ; and for any $\ell_0 \in L_0 \in \mathcal{L}_0^*$ and $\ell_1 \in L_1 \in \mathcal{A}_1^*$ with $\rho_0(\ell_0; L_0) > 0$, define $\rho_1^Z(\ell_1; L_1 | L_0, \ell_0) := \rho_1(\ell_1; L_1 | (L_0, L_1), (\ell_0, L_1))$. Note that ρ^Z is well-defined given the assumption that ρ admits a BEU representation.

6.1 Consumption Persistence and Taste Persistence

One natural notion of choice persistence (e.g., Keane, 1997) is that the agent is more likely to choose a particular consumption option today if she chose this option yesterday compared with the scenario in which she chose some other option yesterday. To formalize this notion in our framework, we additionally impose the restriction that today's menu does not contain any new consumption options relative to yesterday's menu.

Axiom 9 (Consumption persistence). For any $L_0 \in \mathcal{L}_0^*$ and $L_1 \in \mathcal{A}_1^*$ with $L_1 \subseteq L_0$,

$$\rho_0^Z(\ell; L_0), \rho_0^Z(\ell'; L_0) > 0 \implies \rho_1^Z(\ell; L_1 | L_0, \ell) \geq \rho_1^Z(\ell'; L_1 | L_0, \ell').$$

Proposition 4 shows that consumption persistence is equivalent to the following notion of *taste persistence*: If yesterday's felicity was (ordinally equivalent to) u , today's felicity is more likely to remain in any convex neighborhood D of u compared with the scenario where yesterday's felicity was some other u' . To state this formally, given any set $D \subseteq \mathbb{R}^Z$ of felicities, let $[D] := \{w \in \mathbb{R}^Z : w \approx v \text{ for some } v \in D\}$.

Proposition 4. Suppose ρ admits a BEU representation $(\Omega, \mathcal{F}^*, \mu, (\mathcal{F}_t, U_t, W_t, u_t))$ and Condition 1 holds. Then ρ^Z satisfies Axiom 9 if and only if for any $u, u' \in \mathbb{R}^Z$ with $\mu(u_0 \approx u)$, $\mu(u_0 \approx u') > 0$ and any convex $D \subseteq \mathbb{R}^Z$ with $u \in D$, we have $\mu(u_1 \in [D] | u_0 \approx u) \geq \mu(u_1 \in [D] | u_0 \approx u')$.

All proofs for Section 6 appear in Supplementary Appendix H. In addition to absolute consumption persistence, we can also compare two choice rules ρ and $\hat{\rho}$ in terms of their consumption persistence:

Definition 7. ρ^Z features *more consumption persistence* than $\hat{\rho}^Z$ if $\rho_0^Z = \hat{\rho}_0^Z$ and for any $L_0 \in \mathcal{L}_0^*$ and $L_1 \in \mathcal{A}_1^*$ with $L_1 \subseteq L_0$,

$$\rho_0^Z(\ell; L_0) > 0 \implies \rho_1^Z(\ell; L_1 | L_0, \ell) \geq \hat{\rho}_1^Z(\ell; L_1 | L_0, \ell).$$

Proposition 5 shows that more consumption persistence corresponds to more taste persistence, in the sense that today's felicity is always more likely to remain in a convex neighborhood

of yesterday’s felicity. For this to be the case, we require that there exists a *joint uniformly ranked pair* of consumption lotteries $\bar{\ell}, \underline{\ell} \in \Delta(Z)$ that satisfy Condition 1 for both ρ and $\hat{\rho}$.

Proposition 5. Suppose that ρ and $\hat{\rho}$ admit BEU representations $(\Omega, \mathcal{F}^*, \mu, (\mathcal{F}_t, U_t, W_t, u_t)), (\hat{\Omega}, \hat{\mathcal{F}}^*, \hat{\mu}, (\hat{\mathcal{F}}_t, \hat{U}_t, \hat{W}_t, \hat{u}_t))$ and there exists a joint uniformly ranked pair. Then ρ^Z features more consumption persistence than $\hat{\rho}^Z$ if and only if for any $u \in \mathbb{R}^Z$ and convex $D \subseteq \mathbb{R}^Z$ with $u \in D$ and $\mu(u_0 \approx u) > 0$, we have $\mu(u_0 \approx u) = \hat{\mu}(\hat{u}_0 \approx u)$ and $\mu(u_1 \in [D] \mid u_0 \approx u) \geq \hat{\mu}(\hat{u}_1 \in [D] \mid \hat{u}_0 \approx u)$.

The following example applies Propositions 4 and 5 to the special case of BEU in which felicities u_t follow a finite stationary Markov chain. We show that in this setting our general notion of consumption persistence entails sharp restrictions on the agent’s felicity process: Axiom 9 holds if and only if the Markov chain is a *renewal process*, where a single parameter α captures the extent of the agent’s taste persistence. Moreover, behavior in this case is equivalent to Jeuland’s (1979) classical notion of “brand loyalty,” whereby repeated choices from any fixed menu follow a renewal process.

Example 5 (Markov evolving utility). Let $\mathcal{U} = \{u^1, \dots, u^m\}$ denote a finite set of possible felicities, where $u^i \not\approx u^j$ for any $i \neq j$ and there exist $\bar{\ell}, \underline{\ell} \in \Delta(Z)$ such that $u^i(\bar{\ell}) > u^i(\underline{\ell})$ for all i . Let M be an irreducible transition matrix, where M_{ij} denotes the probability that period $t + 1$ felicity is u^j conditional on period t felicity being u^i . Assume that the initial distribution $\nu \in \Delta(\mathcal{U})$ has full support and equals the stationary distribution. Any such Markov chain (\mathcal{U}, M, ν) generates a (*stationary*) *Markov evolving utility* representation.⁵⁸ We impose a regularity condition, *non-collinearity*, on felicities in \mathcal{U} , whereby for any i, j, k, l with $i \notin \{j, k, l\}$, we have $u^i \notin [\text{co}\{u^j, u^k, u^l\}]$; this is generically satisfied if the outcome space is rich enough relative to the number of felicities.

Corollary 1. Suppose that ρ has a Markov evolving utility representation (\mathcal{U}, M, ν) satisfying non-collinearity. Then the following are equivalent:

- (i). ρ^Z satisfies Axiom 9;
- (ii). (\mathcal{U}, M, ν) is a renewal process: there exists $\alpha \in [0, 1)$ such that $M_{ii} = \alpha + (1 - \alpha)\nu(u^i)$ and $M_{ij} = (1 - \alpha)\nu(u^j)$ for all $i \neq j$;
- (iii). choices from fixed menus follow a renewal process: for any $L = \{\ell^1, \dots, \ell^n\} \in \mathcal{L}_0^*$, there exists $\beta \in [0, 1)$ such that $\rho_1^Z(\ell^i; L \mid L, \ell^i) = \beta + (1 - \beta)\rho_0^Z(\ell^i; L)$ and $\rho_1^Z(\ell^j; L \mid L, \ell^i) = (1 - \beta)\rho_0^Z(\ell^j; L)$ for any $i \neq j$.

⁵⁸Of course, in the two-period setting any BEU representation is Markov (though not necessarily stationary and full support). In Supplementary Appendix I.2, we characterize stationary Markov evolving utility for arbitrary horizon T . Moreover, as evident from the proof, Corollary 1 remains valid for arbitrary T .

In addition, if ρ and $\hat{\rho}$ admit renewal process representations as in Corollary 1, more consumption persistence corresponds to a higher taste persistence parameter α and the same stationary felicity distribution ν :

Corollary 2. Suppose that ρ and $\hat{\rho}$ have stationary renewal process representations induced by $(\mathcal{U}, \nu, \alpha)$ and $(\hat{\mathcal{U}}, \hat{\nu}, \hat{\alpha})$ respectively. Then ρ^Z features more consumption persistence than $\hat{\rho}^Z$ if and only if $\alpha \geq \hat{\alpha}$ and there exists a bijection $\phi : \mathcal{U} \rightarrow \hat{\mathcal{U}}$ such that $u \approx \phi(u)$ and $\nu(u) = \hat{\nu}(\phi(u))$ for each $u \in \mathcal{U}$. ▲

6.2 Consumption Inertia and Learning

Another setting where one should expect to observe some form of choice persistence is Bayesian evolving beliefs. Indeed, in this case the agent's choices in both periods 0 and 1 reflect her expectation of the *same* fixed but unknown tastes. However, consumption persistence in the sense of Axiom 9 is neither implied by nor implies BEB. Instead, Proposition 6 shows that BEB entails the following form of *consumption inertia*: If the agent chose ℓ yesterday from a menu that also contained ℓ' and today faces the binary choice between ℓ and ℓ' , then she continues to choose ℓ with positive probability. Moreover, on the domain of atemporal consumption problems, this testable implication fully captures the additional behavioral content of BEB relative to BEU, thus providing an alternative characterization to Theorem 3 on this domain.

Axiom 10 (Consumption inertia). For any $L_0 \in \mathcal{L}_0^*$ and $\ell, \ell' \in L_0$ with $\{\ell, \ell'\} \in \mathcal{A}_1^*$,

$$\rho_0^Z(\ell; L_0) > 0 \implies \rho_1^Z(\ell; \{\ell, \ell'\} | L_0, \ell) > 0.$$

Proposition 6. Suppose that ρ admits a BEU representation and Condition 1 holds. Then ρ^Z satisfies Axiom 10 if and only if ρ^Z admits a BEB representation.⁵⁹

Similar to Axiom 8, the intuition is based on the martingale property of beliefs. This implies that an agent who expects ℓ to be better than ℓ' in period 0 must with positive probability continue to expect this in period 1. The restriction to binary period-1 menus in Axiom 10 is crucial: For instance, an agent who in period 0 is unsure whether her ranking is $\ell' \succ \ell \succ \ell''$ or $\ell'' \succ \ell \succ \ell'$ might choose ℓ over both of the other two options, but upon learning her preferences in period 1 would never choose ℓ from $\{\ell, \ell', \ell''\}$.⁶⁰

⁵⁹That is, there exists $(\Omega, \mu, \mathcal{F}^*, (\mathcal{F}_t))$ and an \mathcal{F}^* -measurable felicity \tilde{u} such that $\rho_0^Z(\ell_0; L_0) = \mu(\ell_0 = \operatorname{argmax}_{L_0} u_0)$ and $\rho_1^Z(\ell_1; L_1 | L_0, \ell_0) = \mu(\ell_1 = \operatorname{argmax}_{L_1} u_1 | \ell_0 = \operatorname{argmax}_{L_0} u_0)$, where $u_t = \mathbb{E}[\tilde{u} | \mathcal{F}_t]$ for $t = 0, 1$.

⁶⁰In an earlier working paper version, we analyzed a stronger form of consumption inertia, whereby $\rho_0^Z(\ell; L_0) > 0$ implies $\rho_1^Z(\ell; L_1 | L_0, \ell) > 0$ for all $L_1 \subseteq L_0$. We showed that this is equivalent to the requirement that $\mu(u_1 \approx u | u_0 \approx u) > 0$ for all u with $\mu(u_0 \approx u) > 0$. See Section 5.2 of Frick, Iijima, and Strzalecki (2017).

7 Discussion

7.1 Related Literature

An extensive literature studies axiomatic characterizations of random utility models in the *static* setting (Barberá and Pattanaik, 1986; Block and Marschak, 1960; Falmagne, 1978; Luce, 1959; McFadden and Richter, 1990). Our approach incorporates as its static building block the elegant axiomatization of Gul and Pesendorfer (2006) and Ahn and Sarver (2013). As a preliminary step, we extend their result to an infinite outcome space, which is needed since the space of continuation problems in the dynamic model is infinite. This contribution is complementary to Ma (2018) who also provides an infinite-outcome generalization of Gul and Pesendorfer (2006). In contrast to our result, he relies on a stronger regularity condition that rules out the possibility of ties (whereas ties necessarily arise when evaluating continuation problems under BEU) and focuses on the case with continuous vNM utilities. Lu (2016b) studies a model with an objective state space where choice is between Anscombe-Aumann acts; by focusing on state-independent utilities, he traces all randomness of choice to random arrival of signals.⁶¹ While this is similar in spirit to our BEB representation, our state space is subjective and utility can be state-dependent. A recent paper by Lu and Saito (2018a) studies period-0 random choice between consumption lottery streams and attributes the randomness in choices to a stochastic discount factor.⁶²

The axiomatic literature on *dynamic* random utility, and more generally dynamic stochastic choice, is relatively sparse. Our choice domain is as in Kreps and Porteus (1978); however, while they study deterministic choice in each period, we focus on random choice in each period. To the best of our knowledge, Fudenberg and Strzalecki (2015) is the first axiomatic study of stochastic choice in general decision trees, but they study only the special case of i.i.d. DDC with logit shocks to actions.⁶³ As we discuss in Section 5, the latter model is a special case of DREU, but is incompatible with BEU because it features very different attitudes toward option value. In addition, because of the i.i.d. assumption, their representation does not give rise to history dependent choice behavior and cannot accommodate phenomena such as learning and choice persistence; likewise, challenges such as limited observability do not arise in their setting. A recent paper by Ke (2018) characterizes a dynamic version of the Luce model, where

⁶¹Lu (2016a) studies an analogous model with state-dependent utilities in an objective state-space setting.

⁶²Other recent contributions by Apesteguia, Ballester, and Lu (2017) and Manzini and Mariotti (2018) respectively study random utility models with linearly ordered choice options and binary support.

⁶³On more limited domains, Gul, Natenzon, and Pesendorfer (2014) study an agent who receives an outcome only once at the end of a decision tree and characterize a generalization of the Luce model. Pennesi (2017), Cerigioni (2017), and Cerreia-Vioglio, Maccheroni, Marinacci, and Rustichini (2017) characterize versions of the Luce model where the analyst observes a sequence of stochastic choices over consumptions. There is also non-axiomatic work studying special cases of our representation where the agent makes a one-time consumption choice at a stopping time, e.g., Fudenberg, Strack, and Strzalecki (2016).

randomness of choices is caused by execution mistakes and there is no serially correlated private information. In contrast to BEU, his model again does not feature positive option value, as larger menus might induce more mistakes. Duraj (2018) builds on our paper and characterizes general dynamic random (expected) utility in an objective state-space setting.

The literature on menu choice (Dekel, Lipman, and Rustichini, 2001; Dekel, Lipman, Rustichini, and Sarver, 2007; Dillenberger, Lleras, Sadowski, and Takeoka, 2014; Kreps, 1979) considers an agent’s deterministic preference over menus (or decision trees) at a hypothetical ex-ante stage where the agent does not receive any information but anticipates receiving information later. An important difference of our approach is that we study the agent’s behavior in actual decision trees, allowing information to arrive in each period and therefore focusing on stochastic choice. We discuss the comparison in more detail in Section 2.2.3. Related papers are Krishna and Sadowski (2014, 2016) who study ex-ante preferences over infinite-horizon decision trees and characterize *stationary* versions of our BEU representation. Another related paper by Ahn and Sarver (2013) studies both ex-ante deterministic preference over menus and ex-post stochastic choice from menus; they show how to connect the analysis of Gul and Pesendorfer (2006) and of Dekel, Lipman, and Rustichini (2001) to obtain better identification properties. An adaptation of their sophistication axiom plays a key role in our characterization of BEU.

Finally, an extensive empirical literature uses specifications of discrete choice models in dynamic contexts.⁶⁴ As we discuss in Section 5, our DREU representation nests the most general DDC model, which in turn nests our Bayesian rational BEU model. However, while BEU features only shocks to consumption, most DDC models (in particular, i.i.d. DDC) introduce more general shocks to actions. We show that the latter form of shocks can lead to violations of Bayesian rationality due to the fact that they mechanically apply to continuation menus in a way that is detached from their continuation value. As we discuss, this observation is complementary to Wilcox (2011) and Apesteguia and Ballester (2018), who highlight modeling issues in *static* discrete choice models. In particular, they show that when i.i.d. utility shocks are added to expected utilities, then the probability of choosing a risky option over a safe option can decrease with respect to a risk aversion parameter in the vNM utility.

7.2 Conclusion

This paper provides the first axiomatic analysis of the general model of dynamic random utility and several of its key special cases. Our central axioms restrict how choices across periods are related, capturing the key new testable implications of the dynamic relative to the static model and facilitating comparisons between different versions of dynamic random utility.

In a “backward-looking” direction, we show that while observed choices under dynamic ran-

⁶⁴For surveys, see Rust (1994) and Aguirregabiria and Mira (2010).

dom utility are typically history-dependent, even the most general version of the model entails two history *independence* conditions: Contraction history dependence rules out certain dynamically “irrational” behavior such as the “mere exposure effect,” while linear history independence provides a conceptual justification for a lottery-based procedure to extrapolate behavior across different decision trees. In addition, special cases such as learning or persistent taste shocks impose further testable restrictions on the nature of history dependence that correspond to well-documented forms of choice persistence. In a “forward-looking” direction, we show that Bayesian rationality restricts utility shocks to apply to instantaneous consumptions (as under BEU), creating a tension with desirable statistical properties such as non-degenerate likelihoods that require additional mechanical shocks to continuation menus (as under general DDC).

Our analysis addresses some technical challenges that may be relevant to other work on stochastic choice: In particular, we propose a solution to the limited observability problem that arises from the fact that in dynamic settings past choices typically restrict future opportunity sets; and we extend Gul and Pesendorfer’s (2006) and Ahn and Sarver’s (2013) characterization of static random expected utility to infinite outcome spaces.

Finally, throughout the paper we have restricted attention to stochastic processes (U_t) of utilities that evolve *exogenously*. Here choice behavior appears history-dependent to the analyst due to the fact that past choices partly reveal the agent’s private information. But from the point of view of the agent, past choices have no effect on today’s behavior. However, in many settings it is natural to allow (U_t) to evolve *endogenously*, as a function of the agent’s past consumption: Two prominent examples are habit formation (e.g., Becker and Murphy, 1988), where consuming a certain good in the past may make the agent like it more in the present; and active learning/experimentation, where the agent’s consumption provides information to her about some payoff-relevant state of the world, as modeled for instance by the multi-armed bandit literature (e.g., Gittins and Jones, 1972; Robbins, 1952). This gives rise to an additional form of history dependence, which we term *consumption dependence*, where past consumption directly shapes the agent’s choices today. Nevertheless, as we showed in the previous working paper version, our main insights extend to settings with consumption dependence.⁶⁵ The key idea is to study an enriched primitive, where a history $\mathfrak{h}^{t-1} = (A_0, p_0, z_0, \dots, A_{t-1}, p_{t-1}, z_{t-1})$ now summarizes not only that in each period $k \leq t-1$ the agent faced menu A_k and chose p_k , but also that the agent’s realized consumption was $z_k \in \text{supp } p_k^Z$. Natural adaptations of our axioms to this setting then characterize generalizations of DREU, BEU, and BEB that allow the evolution of the agent’s utility process U_t to be influenced by her past consumption.

⁶⁵See Section 7 of Frick, Iijima, and Strzalecki (2017). The distinction between (what we term) history dependence and consumption dependence goes back to at least Heckman (1981), who highlights the importance of distinguishing these two phenomena, so as to avoid spuriously attributing a causal role to past consumption when observed behavior could instead be explained through serially correlated exogenous utilities (e.g., persistent taste heterogeneity).

Appendix: Main Proofs

The appendix is structured as follows:

- Section A defines equivalent versions of DREU, BEU, and BEB.
- Sections B–D prove (T -period generalizations of) Theorems 1–3.
- Section E collects together several lemmas that are used throughout Sections B–D.

The supplementary appendix contains the following additional material and is available at https://drive.google.com/open?id=1JlRsyzkpi1-0yNfDYoQ_dc3Se6qLBeDM:

- Section F proves Theorem 0.
- Sections G and H collect together proofs for Sections 5 and 6.
- Section I provides additional results on identification and axioms for Markov evolving utility.
- Section J provides all omitted proofs for Sections A, E, and I.

A Equivalent Representations

Instead of working with probabilities over the grand state space Ω , our proofs of Theorems 1–3 will employ equivalent versions of our representations, called S-based representations, that look at one-step-ahead conditionals.⁶⁶ Section A.1 defines S-based representations. Section A.2 establishes the equivalence between DREU, BEU, and BEB representations and their respective S-based analogs.

A.1 S-based Representations

For any $X \in \{X_0, \dots, X_T\}$, $A \in K(\Delta(X))$, $p \in \Delta(X)$, let $N(A, p) := \{U \in \mathbb{R}^X : p \in M(A, U)\}$ and $N^+(A, p) := \{U \in \mathbb{R}^X : \{p\} = M(A, U)\}$.

Definition 8. A *random expected utility (REU) form* on $X \in \{X_0, \dots, X_T\}$ is a tuple $(S, \mu, \{U_s, \tau_s\}_{s \in S})$ where

- S is a finite state space and μ is a probability measure on S
- for each $s \in S$, $U_s \in \mathbb{R}^X$ is a nonconstant utility over X .
- for each $s \in S$, the tie-breaking rule τ_s is a finitely-additive probability measure on the Borel σ -algebra on \mathbb{R}^X and is *proper*, i.e., $\tau_s(N^+(A, p)) = \tau_s(N(A, p))$ for all A, p .

Given any REU form $(S, \mu, \{U_s, \tau_s\}_{s \in S})$ on X_i and any $s \in S$, $A_i \in \mathcal{A}_i$, and $p_i \in \Delta(X_i)$, define

$$\tau_s(p_i, A_i) := \tau_s(\{w \in \mathbb{R}^{X_i} : p_i \in M(M(A_i, U_s), w)\}).$$

Definition 9. An *S-based DREU representation* of ρ consists of tuples $(S_0, \mu_0, \{U_{s_0}, \tau_{s_0}\}_{s_0 \in S_0})$, $(S_t, \{\mu_t^{s_{t-1}}\}_{s_{t-1} \in S_{t-1}}, \{U_{s_t}, \tau_{s_t}\}_{s_t \in S_t})_{1 \leq t \leq T}$ such that for all $t = 0, \dots, T$, we have:

DREU1: For all $s_{t-1} \in S_{t-1}$, $(S_t, \mu_t^{s_{t-1}}, \{U_{s_t}, \tau_{s_t}\}_{s_t \in S_t})$ is an REU form on X_t such that⁶⁷

⁶⁶These are dynamic analogs of the static GP representations in Ahn and Sarver (2013).

⁶⁷For $t = 0$, we abuse notation by letting $\mu_t^{s_{t-1}}$ denote μ_0 for all s_{t-1} .

- (a) $U_{s_t} \not\approx U_{s'_t}$ for any distinct pair $s_t, s'_t \in \text{supp}(\mu_t^{s_{t-1}})$;
- (b) $\text{supp}(\mu_t^{s_{t-1}}) \cap \text{supp}(\mu_t^{s'_{t-1}}) = \emptyset$ for any distinct pair s_{t-1}, s'_{t-1} at $t \geq 1$;
- (c) $\bigcup_{s_{t-1} \in S_{t-1}} \text{supp} \mu_t^{s_{t-1}} = S_t$.

DREU2: For all p_t, A_t , and $h^{t-1} = (A_0, p_0, A_1, p_1, \dots, A_{t-1}, p_{t-1}) \in \mathcal{H}_{t-1}(A_t)$,⁶⁸

$$\rho_t(p_t, A_t | h^{t-1}) = \frac{\sum_{(s_0, \dots, s_t) \in S_0 \times \dots \times S_t} \prod_{k=0}^t \mu_k^{s_{k-1}}(s_k) \tau_{s_k}(p_k, A_k)}{\sum_{(s_0, \dots, s_{t-1}) \in S_0 \times \dots \times S_{t-1}} \prod_{k=0}^{t-1} \mu_k^{s_{k-1}}(s_k) \tau_{s_k}(p_k, A_k)}.$$

An S -based BEU representation of ρ is an S -based DREU representation such that for all $t = 0, \dots, T$, we additionally have:

BEU: For all $s_t \in S_t$, there exists $u_{s_t} \in \mathbb{R}^Z$ such that for all $z_t \in Z, A_{t+1} \in \mathcal{A}_{t+1}$, we have

$$U_{s_t}(z_t, A_{t+1}) = u_{s_t}(z_t) + V_{s_t}(A_{t+1}),$$

where $V_{s_t}(A_{t+1}) := \sum_{s_{t+1}} \mu_{t+1}^{s_t}(s_{t+1}) \max_{p_{t+1} \in A_{t+1}} U_{s_{t+1}}(p_{t+1})$ for $t \leq T-1$ and $V_{s_T} \equiv 0$.

An S -based BEB representation is an S -based BEU representation such that additionally:

BEB: For all $t = 0, \dots, T-1$ and $s_t \in S_t$, there exists $\delta_{s_t} > 0$ such that

$$u_{s_t} = \frac{1}{\delta_{s_t}} \sum_{s_{t+1}} \mu_{t+1}^{s_t}(s_{t+1}) u_{s_{t+1}}.$$

A.2 Equivalence Result

Proposition A.1. Let ρ be a dynamic stochastic choice rule.

- (i). ρ admits a DREU representation if and only if ρ admits an S -based DREU representation.
- (ii). ρ admits a BEU representation if and only if ρ admits an S -based BEU representation.
- (iii). ρ admits a BEB representation if and only if ρ admits an S -based BEB representation.

Proof. See Supplementary Appendix J.1. ■

B Proof of Theorem 1

Instead of establishing the two-period characterization in Theorem 1, this section establishes the characterization of DREU under an arbitrary horizon T . Section B.1 presents T -period generalizations of the axioms from Section 3. Section B.2 introduces important terminology regarding the relationship between states and histories that is used throughout the proofs of Theorems 1–3. Sections B.3 and B.4 then establish sufficiency and necessity directions of the DREU characterization.

⁶⁸For $t = 0$, we again abuse notation by letting $\rho_t(\cdot | h^{t-1})$ denote $\rho_0(\cdot)$ for all h^{t-1} .

B.1 Characterization of DREU for Arbitrary T

For general T , DREU is characterized by straightforward generalizations of Axioms 1–4 from Section 3. We first present the T -period generalizations of Contraction History Independence and Linear History Independence.

Given $h^{t-1} = (A_0, p_0, \dots, A_{t-1}, p_{t-1}) \in \mathcal{H}_{t-1}$, let $(h_{-k}^{t-1}, (A'_k, p'_k))$ denote the sequence of the form $(A_0, p_0, \dots, A'_k, p'_k, \dots, A_{t-1}, p_{t-1})$.⁶⁹ We say that $g^{t-1} \in \mathcal{H}_{t-1}$ is *contraction equivalent* to h^{t-1} if for some k , we have $g^{t-1} = (h_{-k}^{t-1}, (B_k, p_k))$, where $A_k \subseteq B_k$ and $\rho_k(p_k, A_k | h^{k-1}) = \rho_k(p_k, B_k | h^{k-1})$.⁷⁰ That is, g^{t-1} and h^{t-1} differ only in period k , where under g^{t-1} , the agent chooses lottery p_k from menu B_k , while under h^{t-1} , she chooses the same lottery p_k from the contraction $A_k \subseteq B_k$; moreover, conditional on h^{k-1} , the choice of p_k from A_k and the choice of p_k from B_k occur with the same probability. Generalizing Axiom 1, Axiom B.1 requires that choice behavior be the same after h^{t-1} and g^{t-1} :

Axiom B.1 (Contraction History Independence). For all $t \leq T$, if $g^{t-1} \in \mathcal{H}_{t-1}(A_t)$ is contraction equivalent to $h^{t-1} \in \mathcal{H}_{t-1}(A_t)$, then $\rho_t(\cdot, A_t | h^{t-1}) = \rho_t(\cdot, A_t | g^{t-1})$.

We say that a finite set of histories $G^{t-1} \subseteq \mathcal{H}_{t-1}$ is *linearly equivalent* to $h^{t-1} = (A_0, p_0, \dots, A_{t-1}, p_{t-1}) \in \mathcal{H}_{t-1}$ if

$$G^{t-1} = \{(h_{-k}^{t-1}, (\lambda A_k + (1 - \lambda)B_k, \lambda p_k + (1 - \lambda)q_k)) : q_k \in B_k\}$$

for some k , B_k , and $\lambda \in (0, 1]$. That is, G^{t-1} is the collection of histories that differ from h^{t-1} only at period k : Under h^{t-1} , the agent chooses p_k from menu A_k , while G^{t-1} summarizes all possible choices of the form $\lambda p_k + (1 - \lambda)q_k$ from the menu $\lambda A_k + (1 - \lambda)B_k$. Generalizing Axiom 2, Axiom B.2 requires period- t choice behavior following the set of histories G^{t-1} to be the same as conditional on h^{t-1} . To state this formally, define the choice distribution from A_t following $G^{t-1} \subseteq \mathcal{H}_{t-1}(A_t)$,

$$\rho_t(\cdot, A_t | G^{t-1}) := \sum_{g^{t-1} \in G^{t-1}} \rho_t(\cdot, A_t | g^{t-1}) \frac{\rho(g^{t-1})}{\sum_{f^{t-1} \in G^{t-1}} \rho(f^{t-1})},$$

to be the weighted average of all choice distributions $\rho_t(\cdot, A_t | g^{t-1})$ following histories in G^{t-1} , where for each $g^{t-1} = (\hat{A}_0, \hat{p}_0, \dots, \hat{A}_{t-1}, \hat{p}_{t-1})$ its weight $\rho(g^{t-1}) := \prod_{k=0}^{t-1} \rho_k(\hat{p}_k, \hat{A}_k | g^{k-1})$ corresponds to the probability of the sequence of choices summarized by g^{t-1} .⁷¹

Axiom B.2 (Linear History Independence). For all $t \leq T$, if $G^{t-1} \subseteq \mathcal{H}_{t-1}(A_t)$ is linearly equivalent to $h^{t-1} \in \mathcal{H}_{t-1}(A_t)$, then $\rho_t(\cdot, A_t | h^{t-1}) = \rho_t(\cdot, A_t | G^{t-1})$.

Next, we generalize the procedure for overcoming the limited observability problem following arbitrary histories h^{t-1} . To do so, given any menu A_t and history h^{t-1} , consider a degenerate choice sequence $d^{t-1} = (\{q_0\}, q_0, \dots, \{q_{t-1}\}, q_{t-1})$ such that $A_t \in \text{supp } q_{t-1}^A$ and replace $h^{t-1} = (A_0, p_0, \dots, A_{t-1}, p_{t-1})$ with $g^{t-1} := \lambda h^{t-1} + (1 - \lambda)d^{t-1}$ where⁷² at every period $k \leq t - 1$, the agent faces menu $\lambda A_k + (1 - \lambda)\{q_k\}$ and chooses lottery $\lambda p_k + (1 - \lambda)q_k$. Under expected utility

⁶⁹In general this is not a history, but it is if $A'_k \in \text{supp } p_{k-1}^A$ and $A_{k+1} \in \text{supp } p'_k^A$ and $\rho_k(p'_k, A'_k | h^{k-1}) > 0$.

⁷⁰This induces an equivalence relation on \mathcal{H}_{t-1} by taking the symmetric and transitive closure.

⁷¹Note that $\rho(g^{t-1})$ does not keep track of the probabilities $\hat{p}_k^A(\hat{A}_{k+1})$, since these pertain to exogenous randomization and do not reveal any private information.

⁷²In order for $\lambda h^{t-1} + (1 - \lambda)d^{t-1} := (\lambda A_k + (1 - \lambda)\{q_k\}, \lambda p_k + (1 - \lambda)q_k)_{k=0}^{t-1}$ to be a well-defined history, it suffices that $\lambda A_k + (1 - \lambda)\{q_k\} \in \text{supp } q_{k-1}^A$ for all $k = 1, \dots, t - 1$. This can be ensured by appropriately choosing each q_k , working backwards from period $t - 1$.

maximization, g^{t-1} reveals the same information about the agent as h^{t-1} . Thus, we define choices from A_t following h^{t-1} by extrapolating from choices following g^{t-1} .

Define the set of *degenerate period-($t-1$) histories* by $\mathcal{D}_{t-1} := \{d^{t-1} \in \mathcal{H}_{t-1} : d^{t-1} = (\{q_k\}, q_k)_{k=0}^{t-1} \text{ where } q_k \in \Delta(X_k) \forall k \leq t-1\}$.

Definition 10. For any $t \geq 1$, $A_t \in \mathcal{A}_t$, and $h^{t-1} \in \mathcal{H}_{t-1}$, define

$$\rho_t^{h^{t-1}}(\cdot; A_t) := \rho_t(\cdot; A_t | \lambda h^{t-1} + (1-\lambda)d^{t-1}). \quad (10)$$

for some $\lambda \in (0, 1]$ and $d^{t-1} \in \mathcal{D}_{t-1}$ such that $\lambda h^{t-1} + (1-\lambda)d^{t-1} \in \mathcal{H}_{t-1}(A_t)$.

It follows from Axiom B.2 (Linear History Independence) that $\rho_t^{h^{t-1}}(\cdot; A_t)$ is well-defined: Lemma E.4 shows that the RHS of (10) does not depend on the specific choice of λ and d^{t-1} . Moreover, $\rho_t^{h^{t-1}}(\cdot; A_t)$ coincides with $\rho_t(\cdot; A_t | h^{t-1})$ whenever $h^{t-1} \in \mathcal{H}_{t-1}(A_t)$. In the following, we do not distinguish between the extended and nonextended version of ρ_t and use $\rho_t(\cdot; A_t | h^{t-1})$ to denote both.

Generalizing Axiom 3, we now impose the static REU conditions on each extended choice distribution $\rho_t(\cdot | h^{t-1})$:

Axiom B.3 (History-dependent REU). For all $t \leq T$ and h^{t-1} , $\rho_t(\cdot | h^{t-1})$ satisfies Axiom 0.⁷³

Finally, we state the T -period generalization of Axiom 4 (History Continuity). For this, we first define T -period analogs of menus and histories without ties:

Definition 11. For any $0 \leq t \leq T$ and $h^{t-1} \in \mathcal{H}_{t-1}$, the set of period- t *menus without ties* conditional on history h^{t-1} is denoted $\mathcal{A}_t^*(h^{t-1})$ ⁷⁴ and consists of all $A_t \in \mathcal{A}_t$ such that for any $p_t \in A_t$ and any sequences $p_t^n \rightarrow^m p_t$ and $B_t^n \rightarrow^m A_t \setminus \{p_t\}$, we have

$$\lim_{n \rightarrow \infty} \rho_t(p_t^n, B_t^n \cup \{p_t^n\} | h^{t-1}) = \rho_t(p_t, A_t | h^{t-1}).$$

For $t = 0$, we write $\mathcal{A}_0^* := \mathcal{A}_0^*(h^{t-1})$. The set of period t *histories without ties* is $\mathcal{H}_t^* := \{h^t = (A_0, p_0, \dots, A_{t-1}, p_{t-1}) \in \mathcal{H}_t : A_k \in \mathcal{A}_k^*(h^{k-1}) \text{ for all } k \leq t\}$.

We say that $h^{t,n} \rightarrow^m h^t$ if $h^{t,n} = (A_0^n, p_0^n, \dots, A_t^n, p_t^n)$ and $h^t = (A_0, p_0, \dots, A_t, p_t)$ satisfy $A_k^n \rightarrow^m A_k$ and $p_k^n \rightarrow^m p_k$ for each k .

Axiom B.4 (History Continuity). For all $t \leq T-1$, A_{t+1} , p_{t+1} , and h^t ,

$$\rho_{t+1}(p_{t+1}; A_{t+1} | h^t) \in \text{co}\{\lim_n \rho_{t+1}(p_{t+1}; A_{t+1} | h^{t,n}) : h^{t,n} \rightarrow^m h^t, h^{t,n} \in \mathcal{H}_t^*\}.$$

Generalizing Theorem 1, we have the following representation theorem:

Theorem B.1. The dynamic stochastic choice rule ρ satisfies Axioms B.1–B.4 if and only if ρ admits a DREU representation.

⁷³Lemma E.1 verifies that X_t is a separable metric space. Then Mixture Continuity and Finiteness make use of the same convergence notions as defined following Axiom 0.

⁷⁴Note that $\mathcal{A}_t^*(h^{t-1}) \not\subseteq \mathcal{A}_t(h^{t-1})$ because the first set contains all menus without ties (we use history h^{t-1} here only to determine where ties could occur) while the second set contains only menus that occur with positive probability after history h^{t-1} —typically very few menus.

B.2 Relationship between Histories and States

Throughout the proofs of Theorems B.1–D.1 we will make use of the following terminology concerning the relationship between histories and states. Fix any $t \in \{0, \dots, T\}$. Suppose that $(S_{t'}, \{\mu_{s_{t'}}^{s_{t'}-1}\}_{s_{t'}-1 \in S_{t'}-1}, \{U_{s_{t'}}, \tau_{s_{t'}}\}_{s_{t'} \in S_{t'}})$ satisfy DREU1 and DREU2 from Definition 9 for each $t' \leq t$.

Fix any state $s_t^* \in S_t$. We let $\text{pred}(s_t^*)$ denote the unique *predecessor sequence* $(s_0^*, \dots, s_{t-1}^*) \in S_0 \times \dots \times S_{t-1}$, given by assumptions DREU1 (b) and (c), such that $s_{k+1}^* \in \text{supp}(\mu_{s_{k+1}^*}^{s_k^*})$ for each $k = 0, \dots, t-1$. Given any history $h^t = (A_0, p_0, \dots, A_t, p_t)$, we say that s_t^* is *consistent* with h^t if $\prod_{k=0}^t \tau_{s_k^*}(p_k, A_k) > 0$.

For any $k = 0, \dots, t$, $s_k \in S_k$, $p_0 \in A_0 \in \mathcal{A}_0$, and $p_{k+1} \in A_{k+1} \in \mathcal{A}_{k+1}$, let

$$\begin{aligned} \mathcal{U}_{s_k}(A_{k+1}, p_{k+1}) &:= \{U_{s_{k+1}} : s_{k+1} \in \text{supp} \mu_{s_{k+1}}^{s_k} \text{ and } p_{k+1} \in M(A_{k+1}, U_{s_{k+1}})\}; \\ \mathcal{U}_0(A_0, p_0) &:= \{U_{s_0} : s_0 \in S_0 \text{ and } p_0 \in M(A_0, U_{s_0})\}. \end{aligned}$$

A *separating history* for s_t^* is a history $h^t = (B_0, q_0, \dots, B_t, q_t)$ such that $\mathcal{U}_{s_{k-1}^*}(B_k, q_k) = \{U_{s_k^*}\}$ for all $k = 0, \dots, t$ and $h^t \in \mathcal{H}_t^*$, where we abuse notation by letting $\mathcal{U}_{s_{k-1}^*}(B_0, q_0)$ denote $\mathcal{U}_0(B_0, q_0)$. Note that separating histories are required to be histories without ties.

We record the following properties:

Lemma B.1. Fix any $s_t^* \in S_t$ with $\text{pred}(s_t^*) = (s_0^*, \dots, s_{t-1}^*)$. Suppose $h^t = (B_0, q_0, \dots, B_t, q_t)$ satisfies $\mathcal{U}_{s_{k-1}^*}(B_k, q_k) = \{U_{s_k^*}\}$ for all $k = 0, \dots, t$. Then for all $k = 0, \dots, t$, s_k^* is the only state in S_k that is consistent with h^k .

Proof. Fix any $\ell = 0, \dots, t$. First, consider $s'_\ell \in S_\ell \setminus \{s_\ell^*\}$, with $\text{pred}(s'_\ell) = (s'_0, \dots, s'_{\ell-1})$. Let $k \leq \ell$ be smallest such that $s'_k \neq s_k^*$. Then $s'_k \in \text{supp} \mu_{s'_k}^{s_{k-1}^*}$, so $\mathcal{U}_{s_{k-1}^*}(B_k, q_k) = \{U_{s_k^*}\}$ implies that $q_k \notin M(B_k, U_{s'_k})$. Thus, $\tau_{s'_k}(q_k, B_k) = 0$, whence s'_ℓ is not consistent with h^ℓ .

Next, to show that s_ℓ^* is consistent with h^ℓ , note that $\rho_\ell(q_\ell, B_\ell | h^{\ell-1}) > 0$, so DREU2 implies

$$\sum_{(s_0, \dots, s_\ell) \in S_0 \times \dots \times S_\ell} \prod_{k=0}^{\ell} \mu_{s_k}^{s_{k-1}^*}(s_k) \tau_{s_k}(q_k, B_k) > 0. \quad (11)$$

Now, if $(s_0, \dots, s_{\ell-1}) \neq \text{pred}(s_\ell)$, then $\prod_{k=0}^{\ell} \mu_{s_k}^{s_{k-1}^*}(s_k) = 0$. And if $(s_0, \dots, s_{\ell-1}) = \text{pred}(s_\ell)$ but $s_\ell \neq s_\ell^*$, then the first paragraph shows $\prod_{k=0}^{\ell} \tau_{s_k}(q_k, B_k) = 0$. Hence, (11) reduces to $\prod_{k=0}^{\ell} \mu_{s_k}^{s_{k-1}^*}(s_k^*) \tau_{s_k^*}(q_k, B_k) > 0$, whence s_ℓ^* is consistent with h^ℓ . \blacksquare

Lemma B.2. Every $s_t^* \in S_t$ admits a separating history.

Proof. Fix any $s_t^* \in S_t$ with $\text{pred}(s_t^*) = (s_0^*, \dots, s_{t-1}^*)$. By Lemma E.2 and DREU1 (a), there exist menus $B_0 = \{q_0(s_0) : s_0 \in S_0\} \in \mathcal{A}_0$ and $B_k(s_{k-1}) = \{p_k(s_k) : s_k \in \text{supp} \mu_{s_k}^{s_{k-1}^*}\} \in \mathcal{A}_k$ for each $k = 1, \dots, t$ and $s_k \in S_k$ such that $\mathcal{U}_0(B_0, q_0(s_0)) = \{U_{s_0}\}$ for all $s_0 \in S_0$ and $\mathcal{U}_{s_{k-1}^*}(B_k(s_{k-1}), q_k(s_k)) = \{U_{s_k}\}$ for all $s_k \in \text{supp} \mu_{s_k}^{s_{k-1}^*}$. Moreover, we can assume that $B_{k+1}(s_k) \in \text{supp} q_k(s_k)^A$ for all $k = 0, \dots, t-1$ and $s_k \in S_k$, by letting each $q_k(s_k)$ put small enough weight on $(z, B_{k+1}(s_k))$ for some $z \in Z$. Then $h^t := (B_0, q_0(s_0^*), \dots, B_t(s_t^*), q_t(s_t^*(t))) \in \mathcal{H}_t$. Moreover, since $\mathcal{U}_{s_{k-1}^*}(B_k, q_k(s_k^*)) = \{U_{s_k^*}\}$, Lemma B.1 implies that or all for all $k = 0, \dots, t$, s_k^* is the only state consistent with h^k . Additionally, for all $k = 0, \dots, t$ and $s_k \in \text{supp} \mu_{s_k}^{s_{k-1}^*}$, we have $M(B_k(s_{k-1}^*), U_{s_k}) = \{q_k(s_k)\}$ by construction. Hence, by Lemma E.3, we have $B_k(s_{k-1}^*) \in \mathcal{A}_k^*(h^{k-1})$. Thus $h^t \in \mathcal{H}_t^*$, so h^t is a separating history for s_t^* . \blacksquare

B.3 Proof of Theorem B.1: Sufficiency

Suppose ρ satisfies Axioms B.1–B.4. To show that ρ admits a DREU representation, it suffices, by Proposition A.1, to construct an S-based DREU representation for ρ .

We proceed by induction on $t \leq T$. First consider $t = 0$. Since ρ_0 satisfies Axiom B.3 and X_0 is a separable metric space by Lemma E.1, the existence of $(S_0, \mu_0, \{U_{s_0}, \tau_{s_0}\}_{s_0 \in S_0})$ satisfying DREU1 and DREU2 from Definition 9 is immediate from Theorem F.1, which extends Gul and Pesendorfer’s (2006) and Ahn and Sarver’s (2013) characterization result for static S-based REU representations to separable metric spaces and which we prove in Supplementary Appendix F.

Suppose next that $0 \leq t < T$ and that we have constructed $(S_{t'}, \{\mu_{t'}^{s_{t'}-1}\}_{s_{t'}-1 \in S_{t'}-1}, \{U_{s_{t'}}, \tau_{s_{t'}}\}_{s_{t'} \in S_{t'}})$ satisfying DREU1 and DREU2 for each $t' \leq t$. We now construct $(S_{t+1}, \{\mu_{t+1}^{s_t}\}_{s_t \in S_t}, \{U_{s_{t+1}}, \tau_{s_{t+1}}\}_{s_{t+1} \in S_{t+1}})$ satisfying DREU1 and DREU2.

B.3.1 Defining $\rho_{t+1}^{s_t}$ and $(S_{t+1}, \{\mu_{t+1}^{s_t}\}_{s_t \in S_t}, \{U_{s_{t+1}}, \tau_{s_{t+1}}\}_{s_{t+1} \in S_{t+1}})$:

To this end, we first pick an arbitrary separating history $h^t(s_t)$ for each $s_t \in S_t$ (this exists by Lemma B.2) and define

$$\rho_{t+1}^{s_t}(\cdot, A_{t+1}) := \rho_{t+1}(\cdot, A_{t+1} | h^t(s_t))$$

for all $A_{t+1} \in \mathcal{A}_{t+1}$. Note that here $\rho_{t+1}(\cdot, |h^t(s_t))$ is the extended version of $\rho_{t+1}(\cdot | h^t(s_t))$ given in Definition 10; by Axiom B.2 and Lemma E.4, the specific choice of $\lambda \in (0, 1]$ and $d^{t-1} \in \mathcal{D}_{t-1}$ used in the extension procedure does not matter.

By Axiom B.3 and the fact that X_{t+1} is separable metric (Lemma E.1), Theorem F.1 applied to $\rho_{t+1}^{s_t}$ yields an REU form $(S_{t+1}^{s_t}, \mu_{t+1}^{s_t}, \{U_{s_{t+1}}, \tau_{s_{t+1}}\}_{s_{t+1} \in S_{t+1}^{s_t}})$ on X_{t+1} such that $U_{s_{t+1}} \not\approx U_{s'_{t+1}}$ for any distinct pair $s_{t+1}, s'_{t+1} \in S_{t+1}^{s_t}$ and such that

$$\rho_{t+1}^{s_t}(p_{t+1}, A_{t+1}) = \sum_{s_{t+1} \in S_{t+1}^{s_t}} \mu_{t+1}^{s_t}(s_{t+1}) \tau_{s_{t+1}}(p_{t+1}, A_{t+1})$$

for all p_{t+1} and A_{t+1} . Without loss, we can assume that $S_{t+1}^{s_t}$ and $S_{t+1}^{s'_t}$ are disjoint whenever $s_t \neq s'_t$. Set $S_{t+1} := \bigcup_{s_t \in S_t} S_{t+1}^{s_t}$ and extend $\mu_{t+1}^{s_t}$ to a probability measure on S_{t+1} by setting $\mu_{t+1}^{s_t}(s_{t+1}) = 0$ for all $s_{t+1} \in S_{t+1} \setminus S_{t+1}^{s_t}$.

By construction, it is immediate that $(S_{t+1}, \{\mu_{t+1}^{s_t}\}_{s_t \in S_t}, \{U_{s_{t+1}}, \tau_{s_{t+1}}\}_{s_{t+1} \in S_{t+1}})$ thus defined satisfies DREU1 and that

$$\rho_{t+1}^{s_t}(p_{t+1}, A_{t+1}) = \sum_{s_{t+1} \in S_{t+1}} \mu_{t+1}^{s_t}(s_{t+1}) \tau_{s_{t+1}}(p_{t+1}, A_{t+1}) \quad (12)$$

for all p_{t+1} and A_{t+1} . It remains to show that DREU2 is also satisfied.

B.3.2 $\rho_{t+1}^{s_t}$ is Well-Behaved

To this end, Lemma B.3 below first shows that the definition of $\rho_{t+1}^{s_t}$ is well-behaved, in the sense that for any history h^t that can only arise in state s_t , $\rho_{t+1}^{s_t} = \rho_{t+1}(\cdot | h^t)$.

Lemma B.3. Fix any $s_t^* \in S_t$ with $\text{pred}(s_t^*) = (s_0^*, \dots, s_{t-1}^*)$. Suppose $h^t = (A_0, p_0, \dots, A_t, p_t) \in \mathcal{H}_t$ satisfies $\mathcal{U}_{s_{k-1}^*}(A_k, p_k) = \{U_{s_k^*}\}$ for all $k = 0, 1, \dots, t$. Then for any $A_{t+1} \in \mathcal{A}_{t+1}$, $\rho_{t+1}(\cdot, A_{t+1} | h^t) = \rho_{t+1}^{s_t^*}(\cdot, A_{t+1})$.

Proof. Step 1: Let $\tilde{h}^t = (\tilde{A}_0, \tilde{p}_0, \dots, \tilde{A}_t, \tilde{p}_t)$ denote the separating history for s_t^* used to define $\rho_{t+1}^{s_t^*}$. We first prove the Lemma under the assumption that $h^t \in \mathcal{H}_t^*$, i.e., that h^t is itself a separating history for s_t^* .⁷⁵

Pick $(r_0, \dots, r_t) \in \Delta(X_0) \times \dots \times \Delta(X_t)$ such that $A_{t+1} \in \text{supp } r_t^A$ and for all $k = 0, \dots, t-1$,

$$\text{supp}(r_k^A) \supseteq \{B_{k+1}, \tilde{B}_{k+1}, B_{k+1} \cup \tilde{B}_{k+1}\},$$

where $B_\ell := \frac{1}{3}A_\ell + \frac{1}{3}\{\tilde{p}_\ell\} + \frac{1}{3}\{r_\ell\}$ and $\tilde{B}_\ell := \frac{1}{3}\tilde{A}_\ell + \frac{1}{3}\{p_\ell\} + \frac{1}{3}\{r_\ell\}$ for $\ell = 0, \dots, t$. Define $q_\ell := \frac{1}{3}p_\ell + \frac{1}{3}\tilde{p}_\ell + \frac{1}{3}r_\ell$.

Note that since $h^t, \tilde{h}^t \in \mathcal{H}_t^*$ and $\mathcal{U}_{s_{k-1}^*}(A_k, p_k) = \mathcal{U}_{s_{k-1}^*}(\tilde{A}_k, \tilde{p}_k) = \{U_{s_k^*}\}$, Lemma E.3 implies that $M(A_k, U_{s_k^*}) = \{p_k\}$ and $M(\tilde{A}_k, U_{s_k^*}) = \{\tilde{p}_k\}$ for all $k = 0, 1, \dots, t$. By linearity of the U_s , we then also have

$$\begin{aligned} \mathcal{U}_{s_{k-1}^*}(B_k, q_k) &= \mathcal{U}_{s_{k-1}^*}(\tilde{B}_k, q_k) = \mathcal{U}_{s_{k-1}^*}(B_k \cup \tilde{B}_k, q_k) = \{U_{s_k^*}\} \text{ and} \\ M(B_k, U_{s_k^*}) &= M(\tilde{B}_k, U_{s_k^*}) = M(B_k \cup \tilde{B}_k, U_{s_k^*}) = \{q_k\}. \end{aligned}$$

This implies that for all $k = 0, \dots, t$ and $s_k \in \text{supp } \mu_{k-1}^{s_{k-1}^*}$,

$$\tau_{s_k}(q_k, B_k) = \tau_{s_k}(q_k, \tilde{B}_k) = \tau_{s_k}(q_k, B_k \cup \tilde{B}_k) = \begin{cases} 1 & \text{if } s_k = s_k^* \\ 0 & \text{otherwise} \end{cases}$$

By DREU2 of the inductive hypothesis, it follows that for all $k = 0, \dots, t-1$,

$$\begin{aligned} \mu_t^{s_{t-1}^*}(s_t^*) &= \rho_t(q_t, B_t | B_0, q_0, \dots, B_{t-1}, q_{t-1}) = \rho_t(q_t, \tilde{B}_t | \tilde{B}_0, q_0, \dots, \tilde{B}_{t-1}, q_{t-1}) \\ &= \rho_t(q_t, B_t \cup \tilde{B}_t | B_0, q_0, \dots, B_{t-1}, q_{t-1}, B_t \cup \tilde{B}_t, q_t, \dots, B_{t-1} \cup \tilde{B}_{t-1}, q_{t-1}) \\ &= \rho_t(q_t, B_t \cup \tilde{B}_t | \tilde{B}_0, q_0, \dots, \tilde{B}_{t-1}, q_{t-1}, B_t \cup \tilde{B}_t, q_t, \dots, B_{t-1} \cup \tilde{B}_{t-1}, q_{t-1}), \end{aligned}$$

whence repeated application of Axiom B.1 (Contraction History Independence) yields

$$\begin{aligned} \rho_{t+1}(\cdot, A_{t+1} | B_0, q_0, \dots, B_t, q_t) &= \rho_{t+1}(\cdot, A_{t+1} | B_0 \cup \tilde{B}_0, q_0, \dots, B_t \cup \tilde{B}_t, q_t) = \\ &= \rho_{t+1}(\cdot, A_{t+1} | \tilde{B}_0, q_0, \dots, \tilde{B}_t, q_t). \end{aligned} \tag{13}$$

Moreover, by Axiom B.2 (Linear History Independence) and Lemma E.4, we have

$$\begin{aligned} \rho_{t+1}(\cdot, A_{t+1} | h^t) &= \rho_{t+1}(\cdot, A_{t+1} | B_0, q_0, \dots, B_t, q_t) \text{ and} \\ \rho_{t+1}(\cdot, A_{t+1} | \tilde{h}^t) &= \rho_{t+1}(\cdot, A_{t+1} | \tilde{B}_0, q_0, \dots, \tilde{B}_t, q_t). \end{aligned} \tag{14}$$

Combining (13) and (14) we obtain that $\rho_{t+1}(\cdot, A_{t+1} | h^t) = \rho_{t+1}(\cdot, A_{t+1} | \tilde{h}^t) := \rho_{t+1}^{s_t^*}(\cdot, A_{t+1})$. This proves the Lemma for histories $h^t \in \mathcal{H}_t^*$.

Step 2: Now suppose that $h^t \notin \mathcal{H}_t^*$. Take any sequence of histories $h^{t,n} \xrightarrow{m} h^t$ with $h^{t,n} = (A_0^n, p_0^n, \dots, A_t^n, p_t^n) \in \mathcal{H}_t^*$ for each n . Note that such a sequence exists by Axiom B.4 (History Continuity).

We claim that for all large enough n , $\mathcal{U}_{s_{k-1}^*}(A_k^n, p_k^n) = \{U_{s_k^*}\}$ for all $k = 0, \dots, t$. Suppose for a contradiction that we can find a subsequence $(h^{t,n_\ell})_{\ell=1}^\infty$ for which this claim is violated. Note that

⁷⁵Note that $\mathcal{U}_{s_{k-1}^*}(A_k, p_k) = \{U_{s_k^*}\}$ for all $k = 0, 1, \dots, t$ does not by itself imply that h^t is a history without ties.

for all ℓ , $\rho_k(p_k^{n_\ell}, A_k^{n_\ell} | h^{k-1, n_\ell}) > 0$ for all $k = 0, \dots, t$ (by the fact that h^{t, n_ℓ} is a well-defined history). Hence, DREU2 for $k \leq t$ implies that we can find $s'_{t, n_\ell} \in S_t$ with $\text{pred}(s'_{t, n_\ell}) = (s'_{0, n_\ell}, \dots, s'_{t-1, n_\ell})$ and $(s'_{0, n_\ell}, \dots, s'_{t, n_\ell}) \neq (s_0^*, \dots, s_t^*)$ such that $U_{s'_{k, n_\ell}} \in \mathcal{U}_{s'_{k-1, n_\ell}}(A_k^{n_\ell}, p_k^{n_\ell})$ for all $k = 0, \dots, t$. Moreover, since $S_0 \times \dots \times S_t$ is finite, by choosing the subsequence (h^{t, n_ℓ}) appropriately, we can assume that $(s'_{0, n_\ell}, \dots, s'_{t, n_\ell}) = (s'_0, \dots, s'_t) \neq (s_0^*, \dots, s_t^*)$ for all ℓ . Pick the smallest k such that $s'_k \neq s_k^*$ and pick any $q_k \in A_k$. Since $A_k^{n_\ell} \xrightarrow{m} A_k$ we can find $q_k^{n_\ell} \in A_k^{n_\ell}$ with $q_k^{n_\ell} \xrightarrow{m} q_k$. For all ℓ we have $U_{s'_k} \in \mathcal{U}_{s'_{k-1}}(A_k^{n_\ell}, p_k^{n_\ell})$, so $U_{s'_k}(p_k^{n_\ell}) \geq U_{s'_k}(q_k^{n_\ell})$, whence $U_{s'_k}(p_k) \geq U_{s'_k}(q_k)$ by linearity of $U_{s'_k}$. Moreover, by choice of k , $s'_k \in \text{supp } \mu_{k-1}^{s'_{k-1}} = \text{supp } \mu_{k-1}^{s_{k-1}^*}$. Thus, $U_{s'_k} \in \mathcal{U}_{s_{k-1}^*}(A_k, p_k) = \{U_{s_k^*}\}$. But $s'_k \neq s_k^*$, so by DREU1 (a) of the inductive hypothesis $U_{s'_k} \not\approx U_{s_k^*}$, a contradiction.

By the previous paragraph, for large enough n , $h^{t, n}$ satisfies the assumption of the Lemma. Since $h^{t, n} \in \mathcal{H}_t^*$, Step 1 then shows that $\rho_{t+1}(p_{t+1}, A_{t+1} | h^{t, n}) = \rho_{t+1}^{s_t^*}(p_{t+1}, A_{t+1})$ for all large enough n and all p_{t+1} . By Axiom B.4 (History Continuity), this implies that for all p_{t+1}

$$\rho_{t+1}(p_{t+1}, A_{t+1} | h^t) \in \text{co}\{\lim_n \rho_{t+1}(p_{t+1}, A_{t+1} | h^{t, n}) : h^{t, n} \xrightarrow{m} h^t, h^{t, n} \in \mathcal{H}_t^*\} = \{\rho_{t+1}^{s_t^*}(p_{t+1}, A_{t+1})\},$$

which completes the proof. \blacksquare

B.3.3 $\rho_{t+1}(\cdot | h^t)$ is a Weighted Average of $\rho_{t+1}^{s_t}$

The next lemma shows that $\rho_{t+1}(\cdot | h^t)$ can be expressed as a weighted average of the state-dependent choice distributions $\rho_{t+1}^{s_t}$, where the weight on each $\rho_{t+1}^{s_t}$ corresponds to the probability of s_t conditional on history h^t .

Lemma B.4. For any $p_{t+1} \in A_{t+1}$ and $h^t = (A_0, p_0, \dots, A_t, p_t) \in \mathcal{H}_t(A_{t+1})$, we have

$$\rho_{t+1}(p_{t+1}, A_{t+1} | h^t) = \frac{\sum_{(s_0, \dots, s_t) \in S_0 \times \dots \times S_t} \prod_{k=0}^t \mu_k^{s_{k-1}}(s_k) \tau_{s_k}(A_k, p_k) \rho_{t+1}^{s_t}(p_{t+1}, A_{t+1})}{\sum_{(s_0, \dots, s_t) \in S_0 \times \dots \times S_t} \prod_{k=0}^t \mu_k^{s_{k-1}}(s_k) \tau_{s_k}(A_k, p_k)}.$$

Proof. Let $\{s_t^1, \dots, s_t^m\}$ denote the set of states in S_t that are consistent with history h^t (as defined in Section B.2). For each j , let $\hat{h}^t(j) = (B_0^j, q_0^j, \dots, B_t^j, q_t^j)$ be a separating history for state s_t^j . We can assume that for each $k = 1, \dots, t$, q_{k-1}^j puts small weight on $(z, \frac{1}{2}A_k + \frac{1}{2}B_k^j)$ for some z , so that $h^t(j) := \frac{1}{2}h^t + \frac{1}{2}\hat{h}^t(j) \in \mathcal{H}_t(A_{t+1})$ for all j .

Note first that for all $j = 1, \dots, m$, we have

$$\rho(h^t(j)) = \prod_{k=0}^t \mu_k^{s_{k-1}^j}(s_k^j) \tau_{s_k^j}(p_k, A_k). \quad (15)$$

Indeed, observe that

$$\begin{aligned} \rho(h^t(j)) &= \prod_{k=0}^t \rho_k\left(\frac{1}{2}p_k + \frac{1}{2}q_k^j, \frac{1}{2}A_k + \frac{1}{2}B_k^j \mid \frac{1}{2}h^{k-1} + \frac{1}{2}\hat{h}^{k-1}(j)\right) \\ &= \sum_{(s_0, \dots, s_t)} \prod_{k=0}^t \mu_k^{s_{k-1}}(s_k) \tau_{s_k}\left(\frac{1}{2}p_k + \frac{1}{2}q_k^j, \frac{1}{2}A_k + \frac{1}{2}B_k^j\right) \\ &= \prod_{k=0}^t \mu_k^{s_{k-1}^j}(s_k^j) \tau_{s_k^j}\left(\frac{1}{2}p_k + \frac{1}{2}q_k^j, \frac{1}{2}A_k + \frac{1}{2}B_k^j\right) = \prod_{k=0}^t \mu_k^{s_{k-1}^j}(s_k^j) \tau_{s_k^j}(p_k, A_k). \end{aligned}$$

The first equality holds by definition. The second equality follows from DREU2 of the inductive hypothesis. For the final two equalities, note that since $\hat{h}^t(j)$ is a separating history for s_t^j , we have for all $k = 0, \dots, t$ that $\mathcal{U}_{s_{k-1}^j}(B_k^j, q_k^j) = \{U_{s_k^j}\}$ with $\{q_k^j\} = M(B_k^j, U_{s_k^j})$ (by Lemma E.3). Also, since s_t^j is consistent with h^t , $\tau_{s_k^j}(p_k, A_k) > 0$ for all $k = 0, \dots, t$. This implies that for every $s_k \in \text{supp } \mu_k^{s_k^{j-1}}$, $\tau_{s_k}(\frac{1}{2}p_k + \frac{1}{2}q_k^j, \frac{1}{2}A_k + \frac{1}{2}B_k) > 0$ if and only if $s_k = s_k^j$, yielding the third equality. It also implies that $M(\frac{1}{2}A_k + \frac{1}{2}B_k^j, U_{s_k^j}) = M(\frac{1}{2}A_k + \frac{1}{2}\{q_k^j\}, U_{s_k^j})$, so that $\tau_{s_k^j}(\frac{1}{2}p_k + \frac{1}{2}q_k^j, \frac{1}{2}A_k + \frac{1}{2}B_k^j) = \tau_{s_k^j}(\frac{1}{2}p_k + \frac{1}{2}q_k^j, \frac{1}{2}A_k + \frac{1}{2}\{q_k^j\}) = \tau_{s_k^j}(p_k, A_k)$, yielding the fourth equality.

Now let $H^t := \{h^t(j) : j = 1, \dots, m\} \subseteq \mathcal{H}_t(A_{t+1})$. Note that by repeated application of Axiom B.2, we have that

$$\rho_{t+1}(p_{t+1}, A_{t+1}|h^t) = \rho_{t+1}(p_{t+1}, A_{t+1}|H^t). \quad (16)$$

Moreover, we have that

$$\begin{aligned} \rho_{t+1}(p_{t+1}, A_{t+1}|H^t) &= \frac{\sum_{j=1}^m \rho(h^t(j)) \rho_{t+1}(p_{t+1}, A_{t+1}|h^t(j))}{\sum_{j=1}^m \rho(h^t(j))} \\ &= \frac{\sum_{j=1}^m \prod_{k=0}^t \mu_k^{s_k^{j-1}}(s_k^j) \tau_{s_k^j}(p_k, A_k) \rho_{t+1}(p_{t+1}, A_{t+1}|h^t(j))}{\sum_{j=1}^m \prod_{k=0}^t \mu_k^{s_k^{j-1}}(s_k^j) \tau_{s_k^j}(p_k, A_k)} \\ &= \frac{\sum_j \prod_{k=0}^t \mu_k^{s_k^{j-1}}(s_k^j) \tau_{s_k^j}(p_k, A_k) \rho_{t+1}^{s_t^j}(p_{t+1}|A_{t+1})}{\sum_j \prod_{k=0}^t \mu_k^{s_k^{j-1}}(s_k^j) \tau_{s_k^j}(p_k, A_k)} \\ &= \frac{\sum_{(s_0, \dots, s_t) \in S_0 \times \dots \times S_t} \prod_{k=0}^t \mu_k^{s_k^{j-1}}(s_k) \tau_{s_k}(A_k, p_k) \rho_{t+1}^{s_t}(p_{t+1}|A_{t+1})}{\sum_{(s_0, \dots, s_t) \in S_0 \times \dots \times S_t} \prod_{k=0}^t \mu_k^{s_k^{j-1}}(s_k) \tau_{s_k}(A_k, p_k)}. \end{aligned} \quad (17)$$

Indeed, the first equality holds by definition of choice conditional on a set of histories. The second equality follows from Equation (15). Note next that since $\hat{h}^t(j)$ is a separating history for s_t^j and s_t^j is consistent with h^t , we have that $\mathcal{U}_{s_k^j}(\frac{1}{2}p_k + \frac{1}{2}q_k^j, \frac{1}{2}A_k + \frac{1}{2}B_k^j) = \{U_{s_k^j}\}$ for each k . Hence, Lemma B.3 implies that $\rho_{t+1}(p_{t+1}, A_{t+1}|h^t(j)) = \rho_{t+1}^{s_t^j}(p_{t+1}, A_{t+1})$, yielding the third equality. Finally, note that if $(s_0, \dots, s_t) \in S_0 \times \dots \times S_t$ with $(s_0, \dots, s_t) \neq (s_0^j, \dots, s_t^j)$ for all j , then either $s_t \notin \{s_t^1, \dots, s_t^m\}$, or $s_t = s_t^j$ for some j but $(s_0, \dots, s_{t-1}) \neq \text{pred}(s_t^j)$. In either case, $\prod_{k=0}^t \mu_k^{s_k^{j-1}}(s_k) \tau_{s_k}(A_k, p_k) = 0$, yielding the final equality. Combining (16) and (17), we obtain the desired conclusion. ■

B.3.4 Completing the Proof

Finally, combining Lemma B.4 with the representation of $\rho_{t+1}^{s_t}$ in (12) yields that for any $h^t = (A_0, p_0, \dots, A_t, p_t) \in \mathcal{H}_t(A_{t+1})$

$$\begin{aligned} &\rho_{t+1}(p_{t+1}, A_{t+1}|h^t) \\ &= \frac{\sum_{(s_0, \dots, s_t) \in S_0 \times \dots \times S_t} \prod_{k=0}^t \mu_k^{s_k^{j-1}}(s_k) \tau_{s_k}(A_k, p_k) \sum_{s_{t+1} \in S_{t+1}} \mu_{t+1}^{s_t}(s_{t+1}) \tau_{s_{t+1}}(p_{t+1}, A_{t+1})}{\sum_{(s_0, \dots, s_t) \in S_0 \times \dots \times S_t} \prod_{k=0}^t \mu_k^{s_k^{j-1}}(s_k) \tau_{s_k}(A_k, p_k)} \\ &= \frac{\sum_{(s_0, \dots, s_t, s_{t+1}) \in S_0 \times \dots \times S_t \times S_{t+1}} \prod_{k=0}^{t+1} \mu_k^{s_k^{j-1}}(s_k) \tau_{s_k}(A_k, p_k)}{\sum_{(s_0, \dots, s_t) \in S_0 \times \dots \times S_t} \prod_{k=0}^t \mu_k^{s_k^{j-1}}(s_k) \tau_{s_k}(A_k, p_k)}. \end{aligned}$$

Thus, $(S_{t+1}, \{\mu_{t+1}^{s_t}\}_{s_t \in S_t}, \{U_{s_{t+1}}, \tau_{s_{t+1}}\}_{s_{t+1} \in S_{t+1}})$ also satisfies requirement DREU2, completing the proof.

B.4 Proof of Theorem B.1: Necessity

Suppose ρ admits a DREU representation. By Proposition A.1, ρ admits an S-based DREU representation. By Lemma E.5, for each t and $h^t \in \mathcal{H}_t$, the (static) stochastic choice rule $\rho_t(\cdot|h^t) : \mathcal{A}_t \rightarrow \Delta(\Delta(X_t))$ given by the extended version of ρ from Definition 10 also satisfies DREU2. In other words, $\rho_t(\cdot|h^t)$ admits an S-based REU representation (see Definition 12). Thus, Theorem F.1 implies that Axiom B.3 holds. It remains to verify that Axioms B.1, B.2, and B.4 are satisfied.

Claim 1. ρ satisfies Axiom B.1 (Contraction History Independence).

Proof. Take any $h^{t-1} = (h_{-k}^{t-1}, (A_k, p_k))$, $\hat{h}^{t-1} = (h_{-k}^{t-1}, (B_k, p_k)) \in \mathcal{H}_{t-1}(A_t)$ such that $B_k \supseteq A_k$ and $\rho_k(p_k; A_k|h^{k-1}) = \rho_k(p_k; B_k|h^{k-1})$. From DREU2 for ρ_k , $\rho_k(p_k; A_k|h^{k-1}) = \rho_k(p_k; B_k|h^{k-1})$ implies that

$$\sum_{(s_0, \dots, s_k)} \prod_{l=0}^{k-1} \mu_l^{s_{l-1}}(s_l) \tau_{s_l}(p_l, A_l) \mu_k^{s_{k-1}}(s_k) \tau_{s_k}(p_k, A_k) = \sum_{(s_0, \dots, s_k)} \prod_{l=0}^{k-1} \mu_l^{s_{l-1}}(s_l) \tau_{s_l}(p_l, A_l) \mu_k^{s_{k-1}}(s_k) \tau_{s_k}(p_k, B_k). \quad (18)$$

Since $B_k \supseteq A_k$ implies $\tau_{s_k}(p_k, A_k) \geq \tau_{s_k}(p_k, B_k)$ for all s_k , the only way for (18) to hold is if $\tau_{s_k}(p_k, A_k) = \tau_{s_k}(p_k, B_k)$ for all s_k consistent with h^k . Thus,

$$\rho_t(p_t; A_t|h^{t-1}) = \frac{\sum_{(s_0, \dots, s_t) \in S_0 \times \dots \times S_t} \prod_{l=0}^t \mu_l^{s_{l-1}}(s_l) \tau_{s_l}(p_l, A_l)}{\sum_{(s_0, \dots, s_{t-1}) \in S_0 \times \dots \times S_{t-1}} \prod_{l=0}^{t-1} \mu_l^{s_{l-1}}(s_l) \tau_{s_l}(p_l, A_l)} = \rho_t(p_t; A_t|\hat{h}^{t-1}),$$

as required. ■

Claim 2. ρ satisfies Axiom B.2 (Linear History Independence).

Proof. Take any A_t , $h^{t-1} = (A_0, p_0, \dots, A_{t-1}, p_{t-1}) \in \mathcal{H}_{t-1}(A_t)$, and $H^{t-1} \subseteq \mathcal{H}_{t-1}(A_t)$ of the form $H^{t-1} = \{(h_{-k}^{t-1}, (\lambda A_k + (1-\lambda)B_k, \lambda p_k + (1-\lambda)q_k)) : q_k \in B_k\}$ for some $k < t$, $\lambda \in (0, 1)$, and $B_k = \{q_k^j : j = 1, \dots, m\} \in \mathcal{A}_k$. Let $\tilde{A}_k := \lambda A_k + (1-\lambda)B_k$, and for each $j = 1, \dots, m$, let $\tilde{p}_k^j := \lambda p_k + (1-\lambda)q_k^j$ and $\tilde{h}^{t-1}(j) := (h_{-k}^{t-1}, (\tilde{A}_k, \tilde{p}_k^j))$.

By DREU2, for all p_t , we have

$$\rho_t(p_t; A_t|h^{t-1}) = \frac{\sum_{(s_0, \dots, s_t)} \prod_{\ell=0}^t \mu_\ell^{s_{\ell-1}}(s_\ell) \tau_{s_\ell}(p_\ell, A_\ell)}{\sum_{(s_0, \dots, s_{t-1})} \prod_{\ell=0}^{t-1} \mu_\ell^{s_{\ell-1}}(s_\ell) \tau_{s_\ell}(p_\ell, A_\ell)}. \quad (19)$$

Moreover, by definition

$$\rho_t(p_t; A_t|H^{t-1}) = \frac{\sum_{j=1}^m \rho(\tilde{h}^{t-1}(j)) \rho_t(p_t; A_t|\tilde{h}^{t-1}(j))}{\sum_{j=1}^m \rho(\tilde{h}^{t-1}(j))},$$

where for each $j = 1, \dots, m$, DREU2 yields

$$\rho_t(p_t; A_t|\tilde{h}^{t-1}(j)) = \frac{\sum_{(s_0, \dots, s_t)} \left(\prod_{\ell=0, \dots, t; \ell \neq k} \mu_\ell^{s_{\ell-1}}(s_\ell) \tau_{s_\ell}(p_\ell, A_\ell) \right) \mu_k^{s_{k-1}}(s_k) \tau_{s_k}(\tilde{p}_k^j, \tilde{A}_k)}{\sum_{(s_0, \dots, s_{t-1})} \left(\prod_{\ell=0, \dots, t-1; \ell \neq k} \mu_\ell^{s_{\ell-1}}(s_\ell) \tau_{s_\ell}(p_\ell, A_\ell) \right) \mu_k^{s_{k-1}}(s_k) \tau_{s_k}(\tilde{p}_k^j, \tilde{A}_k)}.$$

and

$$\begin{aligned} \rho(\tilde{h}^{t-1}(j)) &:= \prod_{\ell=0, \dots, t-1; \ell \neq k} \rho_\ell(p_\ell; A_\ell | \tilde{h}^{\ell-1}) \rho_k(\tilde{p}_k^j; \tilde{A}_k | \tilde{h}^{k-1}) \\ &= \sum_{(s_0, \dots, s_{t-1})} \left(\prod_{\ell=0, \dots, t-1; \ell \neq k} \mu_\ell^{s_\ell-1}(s_\ell) \tau_{s_\ell}(p_\ell, A_\ell) \right) \mu_k^{s_k-1}(s_k) \tau_{s_k}(\tilde{p}_k^j, \tilde{A}_k). \end{aligned}$$

Combining and rearranging, we obtain

$$\rho_t(p_t; A_t | H^{t-1}) = \frac{\sum_{(s_0, \dots, s_t)} \left(\prod_{\ell=0, \dots, t; \ell \neq k} \mu_\ell^{s_\ell-1}(s_\ell) \tau_{s_\ell}(A_\ell, p_\ell) \right) \mu_k^{s_k-1}(s_k) \sum_{j=1}^m \tau_{s_k}(\tilde{p}_k^j, \tilde{A}_k)}{\sum_{(s_0, \dots, s_{t-1})} \left(\prod_{\ell=0, \dots, t-1; \ell \neq k} \mu_\ell^{s_\ell-1}(s_\ell) \tau_{s_\ell}(A_\ell, p_\ell) \right) \mu_k^{s_k-1}(s_k) \sum_{j=1}^m \tau_{s_k}(\tilde{p}_k^j, \tilde{A}_k)}. \quad (20)$$

But observe that for all s_k ,

$$\begin{aligned} \sum_{j=1}^m \tau_{s_k}(\tilde{p}_k^j, \tilde{A}_k) &= \sum_{j=1}^m \tau_{s_k}(\{w \in \mathbb{R}^{X_k} : \tilde{p}_k^j \in M(M(\tilde{A}_k, U_{s_k}), w)\}) \\ &= \sum_{q_k \in B_k} \tau_{s_k}(\{w \in \mathbb{R}^{X_k} : p_k \in M(M(A_k, U_{s_k}), w) \text{ and } q_k \in M(M(B_k, U_{s_k}), w)\}) \\ &= \tau_{s_k}(\{w \in \mathbb{R}^{X_k} : p_k \in M(M(A_k, U_{s_k}), w)\}) \\ &= \tau_{s_k}(p_k, A_k), \end{aligned} \quad (21)$$

where the second equality follows from linearity of the representation, the third equality from the fact that τ_{s_k} is a proper finitely-additive probability measure on \mathbb{R}^{X_k} , and the remaining equalities hold by definition. Combining (19), (20), and (21), we obtain $\rho_t(p_t; A_t | h^{t-1}) = \rho_t(p_t; A_t | H^{t-1})$, as required. \blacksquare

Claim 3. ρ satisfies Axiom B.4 (History Continuity).

Proof. Fix any $A_t, p_t \in A_t$, and $h^{t-1} = (A_0, p_0, \dots, A_{t-1}, p_{t-1}) \in \mathcal{H}^{t-1}$. Let $S_{t-1}(h^{t-1}) \subseteq S_{t-1}$ denote the set of period- $(t-1)$ states that are consistent with h^{t-1} . Define $\rho_t^{s_{t-1}}(p_t; A_t) := \sum_{s_t} \mu_t^{s_t-1}(s_t) \tau_{s_t}(p_t, A_t)$ for each s_{t-1} . By Lemma E.5,

$$\begin{aligned} \rho_t(p_t; A_t | h^{t-1}) &= \frac{\sum_{(s_0, \dots, s_t) \in S_0 \times \dots \times S_t} \prod_{k=0}^t \mu_k^{s_k-1}(s_k) \tau_{s_k}(p_k, A_k)}{\sum_{(s_0, \dots, s_{t-1}) \in S_0 \times \dots \times S_{t-1}} \prod_{k=0}^{t-1} \mu_k^{s_k-1}(s_k) \tau_{s_k}(p_k, A_k)} \\ &= \frac{\sum_{(s_0, \dots, s_{t-1}) \in S_0 \times \dots \times S_{t-1}} \prod_{k=0}^{t-1} \mu_k^{s_k-1}(s_k) \tau_{s_k}(p_k, A_k) \sum_{s_t \in S_t} \mu_t^{s_t-1}(s_t) \tau_{s_t}(p_t, A_t)}{\sum_{(s_0, \dots, s_{t-1}) \in S_0 \times \dots \times S_{t-1}} \prod_{k=0}^{t-1} \mu_k^{s_k-1}(s_k) \tau_{s_k}(p_k, A_k)}. \end{aligned}$$

Hence, $\rho_t(p_t; A_t | h^{t-1}) \in \text{co}\{\rho_t^{s_{t-1}}(p_t; A_t) : s_{t-1} \in S_{t-1}(h^{t-1})\}$. Fix any $s_{t-1}^* \in S_{t-1}(h^{t-1})$. To prove the claim, it is sufficient to show that

$$\rho_t^{s_{t-1}^*}(p_t; A_t) \in \{\lim_n \rho_t(p_t; A_t | h_n^{t-1}) : h_n^{t-1} \xrightarrow{m} h^{t-1}, h_n^{t-1} \in \mathcal{H}_{t-1}^*\}.$$

To this end, let $\text{pred}(s_{t-1}^*) = (s_0^*, \dots, s_{t-2}^*)$ and let $\bar{h}^{t-1} = (B_0, q_0, \dots, B_{t-1}, q_{t-1}) \in \mathcal{H}_{t-1}^*$ be a separating history for s_{t-1}^* . By Lemma E.6, for each $k = 0, \dots, t-1$, we can find sequences $A_k^n \in \mathcal{A}_k^*(\bar{h}^{k-1})$ and $p_k^n \in A_k^n$ such that $A_k^n \xrightarrow{m} A_k$, $p_k^n \xrightarrow{m} p_k$ and $\mathcal{U}_{s_{k-1}^*}(A_k^n, p_k^n) = \{U_{s_k^*}\}$ for all n

and all $k = 0, \dots, t-1$. Working backwards from $k = t-2$, we can inductively replace A_k^n and p_k^n with a mixture putting small weight on (z, A_{k+1}^n) for some z to ensure that $A_{k+1}^n \in \text{supp } p_k^{n,A}$ for all $k \leq t-2$ while maintaining the properties in the previous sentence. Then by construction $h_n^{t-1} := (A_0^n, p_0^n, \dots, A_{t-1}^n, p_{t-1}^n) \in \mathcal{H}_{t-1}^*(A_t)$ and h_n^{t-1} is a separating history for s_{t-1}^* , which by Lemma E.5 implies

$$\begin{aligned} \rho_t(p_t; A_t | h_n^{t-1}) &= \frac{\sum_{s_t \in S_t} \left(\prod_{k=0}^{t-1} \mu_k^{s_k^*} (s_k^*) \tau_{s_k^*}(p_k, A_k) \right) \mu_t^{s_t^*} (s_t) \tau_{s_t}(p_t, A_t)}{\prod_{k=0}^{t-1} \mu_k^{s_k^*} (s_k^*) \tau_{s_k^*}(p_k, A_k)} \\ &= \sum_{s_t} \mu_t^{s_t^*} (s_t) \tau_{s_t}(p_t, A_t) =: \rho_t^{s_t^*}(p_t; A_t) \end{aligned}$$

for each n . Since $h_n^{t-1} \xrightarrow{m} h^{t-1}$, this verifies the desired claim. \blacksquare

C Proof of Theorem 2

Instead of proving the two-period characterization of BEU in Theorem 2, this section establishes a generalization of Theorem 2 for arbitrary horizon T . Section C.1 presents the T -period axioms for BEU. Sections C.2 and C.3 establish sufficiency and necessity of these axioms.

C.1 Characterization of BEU for Arbitrary T

The following three axioms are straightforward T -period generalizations of Axioms 5–7 from Section 4.1:

Axiom C.1 (Separability). For any history h^{t-1} , A_t and $p_t, q_t \notin A_t$ such that $p_t^Z = q_t^Z$, $p_t^A = q_t^A$, and $A_t \cup \{p_t\}, A_t \cup \{q_t\} \in \mathcal{A}_t^*(h^{t-1})$, we have

$$\rho_t(p_t; A_t \cup \{p_t\} | h^{t-1}) = \rho_t(q_t; A_t \cup \{q_t\} | h^{t-1}).$$

For each t , let m_t, m'_t denote typical elements of $\Delta(\mathcal{A}_t)$, and for each m_t , we let $\bar{A}(m_t)$ denote the average menu induced by m_t , i.e., $\bar{A}(m_t) = \sum_{A_t \in \mathcal{A}_t} m_t(A_t) A_t$.

Axiom C.2 (Stochastic DLR). The following hold for all $t \leq T$ and h^{t-1} :

(i). *Preference for Flexibility*: For any A_{t+1}, B_{t+1} such that $A_{t+1} \subseteq B_{t+1}$ and $\{(z, A_{t+1}), (z, B_{t+1})\} \in \mathcal{A}_t^*(h^{t-1})$,

$$\rho_t((z, B_{t+1}); \{(z, A_{t+1}), (z, B_{t+1})\} | h^{t-1}) = 1.$$

(ii). *Reduction of Mixed Menus*: For any A_t and $(z, m_{t+1}), (z, m'_{t+1}) \notin A_t$ such that $\bar{A}(m_{t+1}) = \bar{A}(m'_{t+1})$ and $A_t \cup \{(z, m_{t+1})\}, A_t \cup \{(z, m'_{t+1})\} \in \mathcal{A}_t^*(h^{t-1})$, we have

$$\rho_t((z, m_{t+1}); A_t \cup \{(z, m_{t+1})\} | h^{t-1}) = \rho_t((z, m'_{t+1}); A_t \cup \{(z, m'_{t+1})\} | h^{t-1}).$$

(iii). *Continuity*: $\rho_t(\cdot | h^{t-1}) : \mathcal{A}_t^*(h^{t-1}) \rightarrow \Delta(\Delta(X_t))$ is continuous.

(iv). *Menu Nondegeneracy*: $\{(z, A_{t+1}), (z, B_{t+1})\} \in \mathcal{A}_t^*(h^{t-1})$ for some z, A_{t+1}, B_{t+1} .

Axiom C.3 (Sophistication). For any $t \leq T-1$, $h^t = (h^{t-1}, A_t, p_t) \in \mathcal{H}_t^*$, z , and $A_{t+1} \subseteq B_{t+1} \in \mathcal{A}_{t+1}^*(h^t)$, the following are equivalent:

- (i). $\rho_{t+1}(p_{t+1}; B_{t+1}|h^t) > 0$ for some $p_{t+1} \in B_{t+1} \setminus A_{t+1}$
- (ii). $\liminf_n \rho_t(\frac{1}{2}p_t + \frac{1}{2}(z, B_{t+1}^n); \frac{1}{2}A_t + \frac{1}{2}\{(z, A_{t+1}^n), (z, B_{t+1}^n)\}|h^{t-1}) > 0$ for all $A_{t+1}^n \rightarrow^m A_{t+1}$, $B_{t+1}^n \rightarrow^m B_{t+1}$.

We have the following T -period generalization of Theorem 2:

Theorem C.1. Suppose that ρ admits a DREU representation. Then ρ satisfies Axioms C.1–C.3 if and only if ρ admits a BEU representation.

C.2 Proof of Theorem C.1: Sufficiency

Throughout this section, we assume that ρ admits a DREU representation and satisfies Axioms C.1–C.3. We will show that ρ admits a BEU representation. By Proposition A.1, it is sufficient to construct an S-based BEU representation. Sections C.2.1–C.2.5 accomplish this.

C.2.1 Recursive Construction up to t

The construction proceeds recursively. Suppose that $t \leq T - 1$. Assume that we have obtained $(S_{t'}, \{\mu_{t'}^{s_{t'}-1}\}_{s_{t'}-1 \in S_{t'}-1}, \{U_{s_{t'}}, \tau_{s_{t'}}\}_{s_{t'} \in S_{t'}})$ for each $t' \leq t$ such that DREU1 and DREU2 hold for all $t' \leq t$ and BEU holds for all $t' \leq t - 1$ (see Definition 9 for the statements of these conditions). Note that the base case $t = 0$ is true because of the fact that ρ admits a DREU representation and by Proposition A.1 (the requirement that BEU holds for $t' \leq t - 1$ is vacuous here). To complete the proof, we will construct $(S_{t+1}, \{\mu_{t+1}^{s_t}\}_{s_t \in S_t}, \{U_{s_{t+1}}, \tau_{s_{t+1}}\}_{s_{t+1} \in S_{t+1}})$ such that DREU1 and DREU2 hold for $t' \leq t + 1$ and BEU holds for $t' \leq t$.

C.2.2 Properties of U_{s_t}

The following lemma translates Axioms C.1 (Separability) and C.2 (Stochastic DLR) into properties of U_{s_t} .

Lemma C.1. For any $s_t \in S_t$, there exist functions $u_{s_t} : Z \rightarrow \mathbb{R}$ and $V_{s_t} : \mathcal{A}_{t+1} \rightarrow \mathbb{R}$ with V_{s_t} non-constant such that

- (i). $U_{s_t}(z_t, A_{t+1}) = u_{s_t}(z_t) + V_{s_t}(A_{t+1})$ for all (z_t, A_{t+1})
- (ii). V_{s_t} is continuous
- (iii). V_{s_t} is linear, i.e., $V_{s_t}(\alpha A_{t+1} + (1 - \alpha)B_{t+1}) = \alpha V_{s_t}(A_{t+1}) + (1 - \alpha)V_{s_t}(B_{t+1})$ for all A_{t+1}, B_{t+1} and $\alpha \in (0, 1)$
- (iv). V_{s_t} is monotone, i.e., $V_{s_t}(A_{t+1}) \leq V_{s_t}(B_{t+1})$ for all $A_{t+1} \subseteq B_{t+1}$.

Proof. Fix any $s_t \in S_t$ and its predecessor $s_{t-1} \in S_{t-1}$ (which is uniquely given by $\mu_{s_t}^{s_{t-1}}(s_t) > 0$). Take a separating history h^{t-1} for s_{t-1} , the existence of which is guaranteed by Lemma B.2. Let S denote the support of $\mu_{s_t}^{s_{t-1}}$.

For (i), it suffices, by standard arguments, to show that $U_{s_t}(\frac{1}{2}(x, A_{t+1}) + \frac{1}{2}(y, B_{t+1})) = U_{s_t}(\frac{1}{2}(x, B_{t+1}) + \frac{1}{2}(y, A_{t+1}))$ for all x, y, A_{t+1}, B_{t+1} . To see this, suppose for a contradiction that $U_{s_t}(\frac{1}{2}(x, A_{t+1}) + \frac{1}{2}(y, B_{t+1})) \neq U_{s_t}(\frac{1}{2}(x, B_{t+1}) + \frac{1}{2}(y, A_{t+1}))$. We only consider the case $U_{s_t}(\frac{1}{2}(x, A_{t+1}) + \frac{1}{2}(y, B_{t+1})) > U_{s_t}(\frac{1}{2}(x, B_{t+1}) + \frac{1}{2}(y, A_{t+1}))$ as the other case is analogous. By applying Lemma E.2 to $\{U_s : s \in S\}$, there exists a menu $A_t = \{r_t^s : s \in S\}$ such that for each $s \in S$, r_t^s is the

unique maximizer of U_s in A_t . By Lemma E.3, $A_t \in \mathcal{A}_t^*(h^{t-1})$. Moreover, we can assume that each r_t^s assigns positive probability to (x, A_{t+1}) , (y, B_{t+1}) , and $(x, B_{t+1}), (y, A_{t+1})$, as otherwise we can mix these three options with all lotteries in A_t (using the same weights for each r_t^s) without affecting the construction. Let $r_t := r_t^{s_t}$ denote the maximizer in state s_t . By choosing ε small enough, we can ensure that $p_t := r_t + \varepsilon(x, A_{t+1}) + \varepsilon(y, B_{t+1}) - \varepsilon(x, B_{t+1}) - \varepsilon(y, A_{t+1})$ and $q_t := r_t - \varepsilon(x, A_{t+1}) - \varepsilon(y, B_{t+1}) + \varepsilon(x, B_{t+1}) + \varepsilon(y, A_{t+1})$ are well-defined lotteries. Note that $p_t^A = q_t^A$ and $p_t^Z = q_t^Z$. Moreover, for small enough ε , we can also ensure that

$$U_{s_t}(p_t) > U_{s_t}(r_t) > U_{s_t}(q_t) > \max_{r'_t \in A_t \setminus \{r_t\}} U_s(r'_t)$$

and

$$U_{s'}(r_t^{s'}) > U_{s'}(p_t), U_{s'}(r_t), U_{s'}(q_t)$$

for all $s' \in S$ with $s_t \neq s'$. Hence, $\rho_t(p_t; A_t \cup \{p_t\} | h^{t-1}) = \mu_t^{s_t-1}(s_t) > 0 = \rho_t(q_t, A_t \cup \{q_t\} | h^{t-1})$ and, by Lemma E.3, $A_t \cup \{p_t\}, A_t \cup \{q_t\} \in \mathcal{A}_t^*(h^{t-1})$. But this contradicts Axiom C.1 (Separability).

Thus, there exist functions $u_{s_t} : Z \rightarrow \mathbb{R}$ and $V_{s_t} : \mathcal{A}_{t+1} \rightarrow \mathbb{R}$ such that $U_{s_t}(z_t, A_{t+1}) = u_{s_t}(z_t) + V_{s_t}(A_{t+1})$ for all z_t and A_{t+1} . Moreover, by Axiom C.2-(iv) (Menu Nondegeneracy) and Lemma E.3, there exist A_{t+1}, B_{t+1} and z_t such that $U_{s_t}(z_t, A_{t+1}) \neq U_{s_t}(z_t, B_{t+1})$. Hence, $V_{s_t}(A_{t+1}) \neq V_{s_t}(B_{t+1})$, so that V_{s_t} is non-constant.

For (ii), Axiom C.2-(iii) (Continuity) together with Proposition F.2 ensures that U_{s_t} is continuous. By part (i), this implies that V_{s_t} is continuous.

For (iii), suppose to the contrary that $V_{s_t}(\alpha A_{t+1} + (1-\alpha)B_{t+1}) \neq \alpha V_{s_t}(A_{t+1}) + (1-\alpha)V_{s_t}(B_{t+1})$ for some α, A_{t+1}, B_{t+1} . We only consider the case $V_{s_t}(\alpha A_{t+1} + (1-\alpha)B_{t+1}) > \alpha V_{s_t}(A_{t+1}) + (1-\alpha)V_{s_t}(B_{t+1})$, as the other case is analogous. Note that the collection $\{V_s : s \in S\}$ induces a finite collection of ordinally distinct vNM utilities V^1, \dots, V^k (with $k \leq |S|$) over \mathcal{A}_{t+1} , all of which are non-constant by part (i). Hence, by Lemma E.2, there exists a finite set $M_{t+1} = \{m_{t+1}^i : i = 1, \dots, k\} \subset \Delta(\mathcal{A}_{t+1})$ of lotteries over \mathcal{A}_{t+1} such that each m_{t+1}^i is the unique maximizer of V^i in M_{t+1} . We can assume that each m_{t+1}^i assigns positive probability to menus $\alpha A_{t+1} + (1-\alpha)B_{t+1}$, A_{t+1} , and B_{t+1} , as otherwise we can mix these three options to all lotteries in M_{t+1} (using the same weights for all m_{t+1}^i) without affecting the construction. Let $m_{t+1}^* \in M_{t+1}$ denote the maximizer of V_{s_t} in M_{t+1} .

By choosing ε small enough, we can ensure that $m_{t+1} := m_{t+1}^* + \varepsilon(\alpha A_{t+1} + (1-\alpha)B_{t+1}) - \varepsilon\alpha A_{t+1} - \varepsilon(1-\alpha)B_{t+1}$ and $m'_{t+1} := m_{t+1}^* - \varepsilon(\alpha A_{t+1} + (1-\alpha)B_{t+1}) + \varepsilon\alpha A_{t+1} + \varepsilon(1-\alpha)B_{t+1}$ are well-defined lotteries in $\Delta(\mathcal{A}_{t+1})$. Note that $\bar{A}(m_{t+1}) = \bar{A}(m'_{t+1})$. Moreover, for small enough $\varepsilon > 0$, we can also ensure that

$$V_{s_t}(m_{t+1}) > V_{s_t}(m_{t+1}^*) > V_{s_t}(m'_{t+1}) > \max_{\tilde{m}_{t+1} \in M_{t+1} \setminus \{m_{t+1}^*\}} V_{s_t}(\tilde{m}_{t+1})$$

and

$$\max_{\tilde{m}_{t+1} \in M_{t+1}} V_{s'_t}(\tilde{m}_{t+1}) > V_{s'_t}(m_{t+1}), V_{s'_t}(m_{t+1}^*), V_{s'_t}(m'_{t+1})$$

for all $s'_t \neq s_t$ in S with $V_{s'_t} \not\approx V_{s_t}$. Fix any $z \in Z$ and let $A_t := \{(z, \tilde{m}_{t+1}) : \tilde{m}_{t+1} \in M_{t+1}\}$. Then Lemma E.3 along with the separability of U_s established in part (i) implies that $\rho_t((z, m_{t+1}); A_t \cup \{(z, m_{t+1})\} | h^{t-1}) = \mu_t^{s_t-1}(\{s : V_s \approx V_{s_t}\}) > 0 = \rho_t((z, m'_{t+1}); A_t \cup \{(z, m'_{t+1})\} | h^{t-1})$. Also $A_t \cup \{(z, m_{t+1})\}, A_t \cup \{(z, m'_{t+1})\} \in \mathcal{A}_t^*(h^{t-1})$. But this contradicts Axiom C.2-(ii) (Reduction of Mixed Menus).

For (iv), suppose to the contrary that $V_{s_t}(B_{t+1}) < V_{s_t}(A_{t+1})$ for some $A_{t+1} \subseteq B_{t+1}$. Let $S_+ := \{s \in S : V_s(B_{t+1}) > V_s(A_{t+1})\}$ and $S_- := \{s \in S : V_s(B_{t+1}) < V_s(A_{t+1})\}$. Note that S_- is nonempty as $s_t \in S_-$. For each $s \in S \setminus (S_+ \cup S_-)$ we take a pair of menus A_{t+1}^s, B_{t+1}^s such that $A_{t+1}^s \subseteq B_{t+1}^s$

and $V_s(A_{t+1}^s) \neq V_s(B_{t+1}^s)$.⁷⁶ Define $A_{t+1}^* := \sum_{s \in S \setminus (S_+ \cup S_-)} \varepsilon_s A_{t+1}^s + (1 - \sum_{s \in S \setminus (S_+ \cup S_-)} \varepsilon_s) A_{t+1}$ and $B_{t+1}^* := \sum_{s \in S \setminus (S_+ \cup S_-)} \varepsilon_s B_{t+1}^s + (1 - \sum_{s \in S \setminus (S_+ \cup S_-)} \varepsilon_s) B_{t+1}$, where $(\varepsilon_s) \in (0, 1)^{S \setminus (S_+ \cup S_-)}$ is a vector such that $\sum_{s \in S \setminus (S_+ \cup S_-)} \varepsilon_s < 1$. Note that $A_{t+1}^* \subseteq B_{t+1}^*$ by construction. Moreover, since each V_s is linear by part (iii), we can choose (ε_s) sufficiently small so that $V_s(A_{t+1}^*) > V_s(B_{t+1}^*)$ for every $s \in S_-$ and $V_s(A_{t+1}^*) < V_s(B_{t+1}^*)$ for every $s \in S_+$. In addition, we can pick (ε_s) to ensure that $V_s(A_{t+1}^*) \neq V_s(B_{t+1}^*)$ for all $s \in S \setminus (S_+ \cup S_-)$. Then $\{(z, A_{t+1}^*), (z, B_{t+1}^*)\} \in \mathcal{A}_t^*(h^{t-1})$, by Lemma E.3. Moreover, $\rho_t((z, A_{t+1}^*); \{(z, A_{t+1}^*), (z, B_{t+1}^*)\} | h^{t-1}) \geq \mu_t^{s_{t-1}}(S_-) > 0$. This contradicts Axiom C.2-(i) (Preference for Flexibility). \blacksquare

C.2.3 Construction of Random Utility in Period $t + 1$

Since ρ admits a DREU representation, it admits an S-based DREU representation by Proposition A.1, so in particular we can obtain $(S_{t+1}, \{\mu_{t+1}^{s_t}\}_{s_t \in S_t}, \{\tilde{U}_{s_{t+1}}, \tau_{s_{t+1}}\}_{s_{t+1} \in S_{t+1}})$ satisfying DREU1 and DREU2 at $t + 1$. For any $s_t \in S_t$, define $\rho_{t+1}^{s_t}$ by $\rho_{t+1}^{s_t}(p_{t+1}, A_{t+1}) := \sum_{s_{t+1}} \mu_{t+1}^{s_t}(s_{t+1}) \tau_{s_{t+1}}(p_{t+1}, A_{t+1})$ for all p_{t+1}, A_{t+1} .

C.2.4 Sophistication and Finiteness of Menu Preference

Before completing the representation, we establish two more lemmas. Using Axiom C.3 (Sophistication), the first lemma ensures that for each s_t , $\rho_{t+1}^{s_t}$ and the preference over \mathcal{A}_{t+1} induced by V_{s_t} satisfy Axioms 1 and 2 in Ahn and Sarver (2013).

Lemma C.2. For any $s_t \in S_t$, separating history h^t for s_t , and $A_{t+1} \subseteq B_{t+1} \in \mathcal{A}_{t+1}^*(h^t)$, the following are equivalent:

- (i). $\rho_{t+1}^{s_t}(B_{t+1} \setminus A_{t+1}; B_{t+1}) > 0$.
- (ii). $V_{s_t}(B_{t+1}) > V_{s_t}(A_{t+1})$.

Proof. Pick any separating history $h^t = (A_0, p_0, \dots, A_t, p_t)$ for s_t . Note that $h^t \in \mathcal{H}_t^*$ by definition. By DREU2 at $t + 1$ and Lemma E.5, we have $\rho_{t+1}(B_{t+1} \setminus A_{t+1}; B_{t+1} | h^t) = \rho_{t+1}^{s_t}(B_{t+1} \setminus A_{t+1}; B_{t+1})$. Thus by Axiom C.3 (Sophistication), it suffices to show that $V_{s_t}(B_{t+1}) > V_{s_t}(A_{t+1})$ if and only if point (ii) in Axiom C.3 holds.

To show the “only if” direction, suppose $V_{s_t}(B_{t+1}) > V_{s_t}(A_{t+1})$ and take any sequences $A_{t+1}^n \rightarrow^m A_{t+1}$ and $B_{t+1}^n \rightarrow^m B_{t+1}$. Since convergence in mixture implies convergence under the Hausdorff metric, we have $\lim_n V_{s_t}(A_{t+1}^n) = V_{s_t}(A_{t+1})$ and $\lim_n V_{s_t}(B_{t+1}^n) = V_{s_t}(B_{t+1})$ by continuity of V_{s_t} (Lemma C.1-(ii)). Hence, there is N such that $V_{s_t}(B_{t+1}^n) > V_{s_t}(A_{t+1}^n)$ for all $n \geq N$. Then for all $n \geq N$, the fact that h^t is a separating history for s_t and $M(A_t, U_{s_t}) = \{p_t\}$ (as $h_t \in \mathcal{H}_t^*$) implies that $M(\frac{1}{2}A_t + \frac{1}{2}\{(z, B_{t+1}^n), (z, A_{t+1}^n)\}, U_{s_t}) = \{\frac{1}{2}p_t + \frac{1}{2}(z, B_{t+1}^n)\}$ for all z . Thus, by DREU2 at t and Lemma E.5, we have $\rho_t(\frac{1}{2}p_t + \frac{1}{2}(z, B_{t+1}^n); \frac{1}{2}A_t + \frac{1}{2}\{(z, B_{t+1}^n), (z, A_{t+1}^n)\} | h^{t-1}) = \rho_t(p_t; A_t | h^{t-1}) > 0$ for all $n \geq N$. That is, point (ii) in Axiom C.3 holds.

For the “if” direction, we prove the contrapositive. Suppose that $V_{s_t}(B_{t+1}) \leq V_{s_t}(A_{t+1})$. Note that since V_{s_t} is monotone and non-constant by Lemma C.1, we have $V_{s_t}(B_{t+1}) = V_{s_t}(A_{t+1}) \neq V_{s_t}(C_{t+1})$ for some C_{t+1} . If $V_{s_t}(A_{t+1}) > V_{s_t}(C_{t+1})$ take $A_{t+1}^n = A_{t+1}$ and $B_{t+1}^n = \frac{n-1}{n}B_{t+1} + \frac{1}{n}C_{t+1}$ for each n , and if $V_{s_t}(A_{t+1}) < V_{s_t}(C_{t+1})$ take $B_{t+1}^n = B_{t+1}$ and $A_{t+1}^n = \frac{n-1}{n}A_{t+1} + \frac{1}{n}C_{t+1}$ for each n . In either case, we have $A_{t+1}^n \rightarrow^m A_{t+1}, B_{t+1}^n \rightarrow^m B_{t+1}$ and $V_{s_t}(B_{t+1}^n) < V_{s_t}(A_{t+1}^n)$ for every n by the linearity of V_{s_t} (Lemma C.1). Combining this with the fact that $M(A_t, U_{s_t}) = \{p_t\}$ (since h^t is a separating

⁷⁶Such a pair exists since each V_s is non-constant. Indeed, if such a pair does not exist for some s , then for any pair of menus $\tilde{A}_{t+1} \neq \tilde{B}_{t+1}$, we have $V_s(\tilde{A}_{t+1}) = V_s(\tilde{A}_{t+1} \cup \tilde{B}_{t+1}) = V_s(\tilde{B}_{t+1})$, a contradiction.

history for s_t), we have $M(\frac{1}{2}A_t + \frac{1}{2}\{(z, B_{t+1}^n), (z, A_{t+1}^n)\}, U_{s_t}) = \{\frac{1}{2}p_t + \frac{1}{2}(z, A_{t+1}^n)\}$ for each n . Given this, DREU2 at t and Lemma E.5 yields $\rho_t(\frac{1}{2}p_t + \frac{1}{2}(z, B_{t+1}^n); \frac{1}{2}A_t + \frac{1}{2}\{(z, B_{t+1}^n), (z, A_{t+1}^n)\}|h^{t-1}) = 0$ for all n . That is, point (ii) in Axiom C.3 does not hold. \blacksquare

The next lemma shows that because of Lemma C.2, the finiteness of each $\text{supp}\mu_{t+1}^{s_t}$ is enough to ensure that the preference over \mathcal{A}_{t+1} induced by each V_{s_t} satisfies Axiom DLR 6 (Finiteness) introduced by Ahn and Sarver (2013):

Lemma C.3. For each $s_t \in S_t$, there is $K_{s_t} > 0$ such that for any A_{t+1} , there is $B_{t+1} \subseteq A_{t+1}$ such that $|B_{t+1}| \leq K_{s_t}$ and $V_{s_t}(A_{t+1}) = V_{s_t}(B_{t+1})$.

Proof. Fix any $s_t \in S_t$ and a separating history h^t for s_t . Let $S_{t+1}(s_t) := \text{supp}\mu_{t+1}^{s_t}$. We will show that $K_{s_t} := |S_{t+1}(s_t)|$ is as required.

Step 1: First consider any $B_{t+1} \in \mathcal{A}_{t+1}^*(h^t)$. Then by Lemma E.3, for each $s_{t+1} \in S_{t+1}(s_t)$ we have $|M(B_{t+1}, \tilde{U}_{s_{t+1}})| = 1$. Letting $A_{t+1} := \bigcup_{s_{t+1} \in S_{t+1}(s_t)} M(B_{t+1}, \tilde{U}_{s_{t+1}})$, we then have that $|A_{t+1}| \leq K_{s_t}$ and $\rho_{t+1}^{s_t}(B_{t+1} \setminus A_{t+1}, B_{t+1}) = 0$. By Lemma C.2, this implies that $V_{s_t}(A_{t+1}) = V_{s_t}(B_{t+1})$, as required.

Step 2: Next take any $B_{t+1} \notin \mathcal{A}_{t+1}^*(h^t)$. By Lemma E.6, we can find a sequence $B_{t+1}^n \rightarrow^m B_{t+1}$ with $B_{t+1}^n \in \mathcal{A}_{t+1}^*(h^t)$ for all n . Then by Step 1, we can find $A_{t+1}^n \subseteq B_{t+1}^n$ for all n such that $|A_{t+1}^n| \leq K_{s_t}$ and $V_{s_t}(A_{t+1}^n) = V_{s_t}(B_{t+1}^n)$. By definition of \rightarrow^m , for each $q_{t+1} \in B_{t+1}$, there exists $D_{t+1}(q_{t+1}) \in \mathcal{A}_{t+1}$ and a sequence $\alpha_n(q_{t+1}) \rightarrow 0$ such that $A_{t+1}^n \subseteq \bigcup_{q_{t+1} \in B_{t+1}} \alpha_n(q_{t+1})D_{t+1}(q_{t+1}) + (1 - \alpha_n(q_{t+1}))\{q_{t+1}\}$ for all n . Hence, since $|A_{t+1}^n| \leq K_{s_t}$ for all n , restricting to a subsequence if necessary, there is $A_{t+1} \subseteq B_{t+1}$ such that $A_{t+1}^n \rightarrow^m A_{t+1}$ and such that $|A_{t+1}| \leq K_{s_t}$. Finally, by continuity of V_{s_t} (Lemma C.1 (ii)), we have $V_{s_t}(B_{t+1}) = V_{s_t}(A_{t+1})$, as required. \blacksquare

C.2.5 Completing the Representation

Recall that in Section C.2.3, we have obtained $(S_{t+1}, \{\mu_{t+1}^{s_t}\}_{s_t \in S_t}, \{\tilde{U}_{s_{t+1}}, \tau_{s_{t+1}}\}_{s_{t+1} \in S_{t+1}})$ satisfying DREU1 and DREU2 at $t + 1$. We now show that for each $s_{t+1} \in S_{t+1}$ there exist $\alpha_{s_{t+1}} > 0$ and $\beta_{s_{t+1}} \in \mathbb{R}$ such that after replacing $\tilde{U}_{s_{t+1}}$ with $U_{s_{t+1}} := \alpha_{s_{t+1}}\tilde{U}_{s_{t+1}} + \beta_{s_{t+1}}$, we additionally have that BEU holds at time t .

Fix any s_t and let $S_{t+1}(s_t) := \text{supp}\mu_{t+1}^{s_t}$. Note that by DREU1 at $t + 1$ and since we have defined $\rho_{t+1}^{s_t}$ by $\rho_{t+1}^{s_t}(p_{t+1}, A_{t+1}) := \sum_{s_{t+1} \in S_{t+1}(s_t)} \mu_{t+1}^{s_t}(s_{t+1})\tau_{s_{t+1}}(p_{t+1}, A_{t+1})$ for all p_{t+1} and A_{t+1} , it follows that $(S_{t+1}(s_t), \mu_{t+1}^{s_t}, \{\tilde{U}_{s_{t+1}}, \tau_{s_{t+1}}\}_{s_{t+1} \in S_{t+1}(s_t)})$ is an S -based REU representation of $\rho_{t+1}^{s_t}$ (see Definition 12).

Since all the $U_{s_{t+1}}$ are non-constant and induce different preferences over $\Delta(X_{t+1})$ for distinct $s_{t+1}, s'_{t+1} \in S_{t+1}(s_t)$ and since V_{s_t} is nonconstant by Lemma C.1, we can find a finite set $Y \subseteq X_{t+1}$ such that (i) V_{s_t} is non-constant on $\mathcal{A}_{t+1}(Y) := \{B_{t+1} \in \mathcal{A}_{t+1} : \cup_{p_{t+1} \in B_{t+1}} \text{supp}(p_{t+1}) \subseteq Y\}$; (ii) for each $s_{t+1} \in S_{t+1}(s_t)$, $\tilde{U}_{s_{t+1}}$ is non-constant on Y ; and (iii) for each distinct pair $s_{t+1}, s'_{t+1} \in S_{t+1}(s_t)$, $\tilde{U}_{s_{t+1}} \not\approx \tilde{U}_{s'_{t+1}}$ on Y .

Observe that by Lemmas C.1 and C.3, the preference \succsim_{s_t} on $\mathcal{A}_{t+1}(Y)$ induced by V_{s_t} satisfies Axioms DLR 1–6 (Weak Order, Continuity, Independence, Monotonicity, Nontriviality, Finiteness) in Ahn and Sarver (2013) (henceforth AS), so by Corollary S1 in AS, \succsim_{s_t} admits a DLR representation (see Definition S1 in AS). Moreover, since $\rho_{t+1}^{s_t}$ admits an S -based REU representation (what AS call a GP representation), so does its restriction to $\mathcal{A}_{t+1}(Y)$. Finally, by Lemma C.2, the pair $(\succsim_{s_t}, \rho_{t+1}^{s_t})$ satisfies AS's Axioms 1 and 2 on $\mathcal{A}_{t+1}(Y)$. Thus, by Theorem 1 in AS, we can find a DLR-GP representation of $(\succsim_{s_t}, \rho_{t+1}^{s_t})$ on $\mathcal{A}_{t+1}(Y)$, i.e., an S -based REU representation $(\hat{S}_{t+1}(s_t), \hat{\mu}_{t+1}^{s_t}, \{\hat{U}_{s_{t+1}}, \hat{\tau}_{s_{t+1}}\}_{s_{t+1} \in \hat{S}_{t+1}(s_t)})$ of $\rho_{t+1}^{s_t}$ on $\mathcal{A}_{t+1}(Y)$ such that \succsim_{s_t} restricted to $\mathcal{A}_{t+1}(Y)$ is represented by \hat{V}_{s_t} , where $\hat{V}_{s_t}(A_{t+1}) := \sum_{s_{t+1} \in \hat{S}_{t+1}(s_t)} \hat{\mu}_{t+1}^{s_t}(s_{t+1}) \max_{p_{t+1} \in A_{t+1}} \hat{U}_{s_{t+1}}(p_{t+1})$. Since V_{s_t} also represents \succsim_{s_t} restricted to $\mathcal{A}_{t+1}(Y)$,

standard arguments yield $\hat{\alpha}_{s_t} > 0$ and $\hat{\beta}_{s_t} \in \mathbb{R}$ such that for all $A_{t+1} \in \mathcal{A}_{t+1}(Y)$, we have $V_{s_t}(A_{t+1}) = \hat{\alpha}_{s_t} \hat{V}_{s_t}(A_{t+1}) + \hat{\beta}_{s_t}$, whence

$$V_{s_t}(A_{t+1}) = \sum_{s_{t+1} \in \hat{S}_{t+1}(s_t)} \hat{\mu}_{t+1}^{s_t}(s_{t+1}) \max_{p_{t+1} \in A_{t+1}} U_{s_{t+1}}(p_{t+1}), \quad (22)$$

where $U_{s_{t+1}} = \hat{\alpha}_{s_t} \hat{U}_{s_{t+1}} + \hat{\beta}_{s_t}$. By the uniqueness properties of S -based REU representations (Proposition 4 in AS), $(\hat{S}_{t+1}(s_t), \hat{\mu}_{t+1}^{s_t}, \{U_{s_{t+1}}, \hat{\tau}_{s_{t+1}}\}_{s_{t+1} \in \hat{S}_{t+1}(s_t)})$ still constitutes an S -based REU representation of $\rho_{t+1}^{s_t}$ on $\mathcal{A}_{t+1}(Y)$. Applying Proposition 4 in AS again, since $(S_{t+1}(s_t), \mu_{t+1}^{s_t}, \{\tilde{U}_{s_{t+1}}, \tau_{s_{t+1}}\}_{s_{t+1} \in S_{t+1}(s_t)})$ also represents $\rho_{t+1}^{s_t}$ on $\mathcal{A}_{t+1}(Y)$, we can assume after re-labeling that $S_{t+1}(s_t) = \hat{S}_{t+1}(s_t)$, $\hat{\mu}_{t+1}^{s_t} = \mu_{t+1}^{s_t}$ and that for each $s_{t+1} \in S_{t+1}(s_t)$, there exist constants $\alpha_{s_{t+1}} > 0$ and $\beta_{s_{t+1}} \in \mathbb{R}$ such that

$$U_{s_{t+1}}(x_{t+1}) = \alpha_{s_{t+1}} \tilde{U}_{s_{t+1}}(x_{t+1}) + \beta_{s_{t+1}} \quad (23)$$

for each $x_{t+1} \in Y \subseteq X_{t+1}$. Since $\tilde{U}_{s_{t+1}}$ is defined on X_{t+1} , we can extend $U_{s_{t+1}}$ to the whole space X_{t+1} by (23). Then $U_{s_{t+1}}$ and $\tilde{U}_{s_{t+1}}$ represent the same preference over $\Delta(X_{t+1})$, so since $(S_{t+1}(s_t), \mu_{t+1}^{s_t}, \{\tilde{U}_{s_{t+1}}, \tau_{s_{t+1}}\}_{s_{t+1} \in S_{t+1}(s_t)})$ satisfies DREU1 and DREU2, so does $(S_{t+1}(s_t), \mu_{t+1}^{s_t}, \{U_{s_{t+1}}, \tau_{s_{t+1}}\}_{s_{t+1} \in S_{t+1}(s_t)})$.

It remains to show that (22) holds for all $A_{t+1} \in \mathcal{A}_{t+1}$, so that BEU is satisfied at s_t . To see this, consider any $A_{t+1} \in \mathcal{A}_{t+1}$ and choose a finite set $Y' \subseteq X_{t+1}$ such that $Y \cup \bigcup_{p_{t+1} \in A_{t+1}} \text{supp}(p_{t+1}) \subseteq Y'$. As above, we can again apply Theorem 1 in AS to obtain a DLR-GP representation $(\bar{S}_{t+1}(s_t), \bar{\mu}_{t+1}^{s_t}, \{\bar{U}_{s_{t+1}}, \bar{\tau}_{s_{t+1}}\}_{s_{t+1} \in \bar{S}_{t+1}(s_t)})$ of the pair $(\bar{\succ}_{s_t}, \rho_{t+1}^{s_t})$ restricted to $\mathcal{A}_{t+1}(Y')$. But since this also yields a DLR-GP representation of $(\bar{\succ}_{s_t}, \rho_{t+1}^{s_t})$ restricted to $\mathcal{A}_{t+1}(Y)$, by the uniqueness property of DLR-GP representations (Theorem 2 in AS), we can assume that $\bar{S}_{t+1}(s_t) = S_{t+1}(s_t)$, $\bar{\mu}_{t+1}^{s_t} = \mu_{t+1}^{s_t}$ and that there exists $\bar{\alpha}_{s_t} > 0$ and $\bar{\beta}_{s_{t+1}} \in \mathbb{R}$ such that $\bar{U}_{s_{t+1}} = \bar{\alpha}_{s_t} U_{s_{t+1}} + \bar{\beta}_{s_{t+1}}$ for each $s_{t+1} \in S_{t+1}(s_t)$. Since $\bar{\succ}_{s_t}$ is represented on $\mathcal{A}_{t+1}(Y')$ by $\bar{V}_{s_t}(B_{t+1}) := \sum_{s_{t+1} \in S_{t+1}(s_t)} \mu_{t+1}^{s_t}(s_{t+1}) \max_{p_{t+1} \in B_{t+1}} \bar{U}_{s_{t+1}}(p_{t+1})$ and since $\bar{\alpha}_{s_t}$ depends only on s_t (and not on s_{t+1}), it follows that $\bar{\succ}_{s_t}$ is also represented on $\mathcal{A}_{t+1}(Y')$ by $V'_{s_t}(B_{t+1}) := \sum_{s_{t+1} \in S_{t+1}(s_t)} \mu_{t+1}^{s_t}(s_{t+1}) \max_{p_{t+1} \in B_{t+1}} U_{s_{t+1}}(p_{t+1})$. Thus, the linear functions V_{s_t} and V'_{s_t} represent the same preference on $\mathcal{A}_{t+1}(Y')$ and coincide on $\mathcal{A}_{t+1}(Y)$, so they must also coincide on $\mathcal{A}_{t+1}(Y')$. Thus, (22) holds at A_{t+1} .

This shows that BEU holds at t . Combining this with the inductive hypothesis, it follows that $(S_{t'}, \{\mu_{t'}^{s_{t'-1}}\}_{s_{t'-1} \in S_{t'-1}}, \{U_{s_{t'}}, \tau_{s_{t'}}\}_{s_{t'} \in S_{t'}})$ satisfies DREU1 and DREU2 for all $t' \leq t+1$ and BEU for all $t' \leq t$, as required.

C.3 Proof of Theorem C.1: Necessity

Suppose that ρ admits a BEU representation. Then by Proposition A.1, ρ admits an S -based BEU representation $(S_t, \{\mu_{t-1}^{s_t}\}_{s_{t-1} \in S_{t-1}}, \{U_{s_t}, u_{s_t}, \tau_{s_t}\}_{s_t \in S_t})$.

To show Axiom C.1 (Separability), take any history h^{t-1} , A_t and $p_t, q_t \notin A_t$ such that $p_t^A = q_t^A$, $p_t^Z = q_t^Z$, and $A_t \cup \{p_t\}, A_t \cup \{q_t\} \in \mathcal{A}_t^*(h^{t-1})$. Note that by the representation $U_{s_t}(p_t) = U_{s_t}(q_t)$ for any s_t . Thus $M(A_t \cup \{p_t\}, U_{s_t}) = M(A_t \cup \{q_t\}, U_{s_t})$ for each s_t . Since $A_t \cup \{p_t\}, A_t \cup \{q_t\} \in \mathcal{A}_t^*(h^{t-1})$, this implies $\rho_t(p_t; A_t \cup \{p_t\} | h^{t-1}) = \rho_t(q_t; A_t \cup \{q_t\} | h^{t-1})$. Axiom C.2-(ii) (Reduction of Mixed Menus) is verified in the same manner, because when $\bar{A}(m_{t+1}) = \bar{A}(m'_{t+1})$, then by the representation $U_{s_t}(z, m_{t+1}) = U_{s_t}(z, m'_{t+1})$ for all z and s_t .

To verify Axiom C.2-(i) (Preference for Flexibility), note that when $A_{t+1} \subseteq B_{t+1}$, then by the representation, we have $U_{s_t}(z, A_{t+1}) \leq U_{s_t}(z, B_{t+1})$ for all z and s_t . Moreover, $\{(z, A_{t+1}), (z, B_{t+1})\} \in$

$\mathcal{A}_t^*(h^{t-1})$ guarantees that the inequality is strict for all s_t with the property that $\mu_t^{s_t-1}(s_t) > 0$ for some s_{t-1} that is consistent with history h^{t-1} . This implies $\rho_t((z, A_{t+1}); \{(z, A_{t+1}), (z, B_{t+1})\} | h^{t-1}) = 1$.

Axiom C.2-(iii) (Continuity) holds by Proposition F.2, because for each s_t the function $U_{s_t} : X_t \rightarrow \mathbb{R}$ is continuous by the representation.

To verify Axiom C.2-(iv) (Menu Nondegeneracy), note that by the representation, U_{s_T} is non-constant for every s_T . Then an inductive argument implies that for any z , $t \leq T - 1$ and s_t , $U_{s_t}(z, \cdot) : \mathcal{A}_{t+1} \rightarrow \mathbb{R}$ is also non-constant. Thus, for each s_t , there is a pair of menus such that $U_{s_t}(z, A_{t+1}^{s_t}) \neq U_{s_t}(z, B_{t+1}^{s_t})$. Define $A_{t+1} := \sum_{s_t \in S_t} \alpha_{s_t} A_{t+1}^{s_t}$ and $B_{t+1} := \sum_{s_t \in S_t} \alpha_{s_t} B_{t+1}^{s_t}$ for some vector $(\alpha_{s_t}) \in (0, 1)^{S_t}$ with $\sum_{s_t \in S_t} \alpha_{s_t} = 1$. Since U_{s_t} is linear in continuation menus by the representation, we can choose (α_{s_t}) such that $U_{s_t}(z, A_{t+1}) \neq U_{s_t}(z, B_{t+1})$ for all s_t . By Lemma E.3, this implies $\{(z, A_{t+1}), (z, B_{t+1})\} \in \mathcal{A}_t^*(h^{t-1})$.

Finally, to show Axiom C.3 (Sophistication), take any history $h^t = (A_0, p_0, \dots, A_t, p_t) \in \mathcal{H}_t^*$, z , and $A_{t+1} \subseteq B_{t+1} \in \mathcal{A}_{t+1}^*(h^t)$. Let $S_t^* \subseteq S_t$ denote the set of states that are consistent with h^t . First note that based on Lemmas E.3 and E.5 and the fact that $B_{t+1} \in \mathcal{A}_{t+1}^*(h^t)$, condition (i) in Axiom C.3 is equivalent to the following condition:

$$(i') : \exists s_t^* \in S_t^*, \exists s_{t+1}^* \in \text{supp} \mu_{t+1}^{s_t^*} \quad \text{such that} \quad \max_{B_{t+1}} U_{s_{t+1}^*} > \max_{A_{t+1}} U_{s_{t+1}^*}.$$

Thus, it suffices to show that condition (i') is equivalent to condition (ii) in Axiom C.3.

Suppose first that condition (i') holds. Then by the representation, we have $U_{s_t^*}(z, B_{t+1}) > U_{s_t^*}(z, A_{t+1})$. Take any sequences $A_{t+1}^n \rightarrow^m A_{t+1}$ and $B_{t+1}^n \rightarrow^m B_{t+1}$. Since convergence in mixture implies convergence under the Hausdorff metric and $U_{s_t^*}$ is continuous by the representation, this yields some N such that $U_{s_t^*}(z, B_{t+1}^n) > U_{s_t^*}(z, A_{t+1}^n)$ for all $n \geq N$. Hence, the fact that p_t is the unique maximizer of $U_{s_t^*}$ in A_t (which follows from $h^t \in \mathcal{H}_t^*$) implies that for all $n \geq N$, $\frac{1}{2}p_t + \frac{1}{2}(z, B_{t+1}^n)$ is the unique maximizer of $U_{s_t^*}$ in $\frac{1}{2}A_t + \frac{1}{2}\{(z, B_{t+1}^n), (z, A_{t+1}^n)\}$. This ensures that for all $n \geq N$, $\rho_t(\frac{1}{2}p_t + \frac{1}{2}(z, B_{t+1}^n); \frac{1}{2}A_t + \frac{1}{2}\{(z, B_{t+1}^n), (z, A_{t+1}^n)\} | h^{t-1})$ is strictly positive, as it is greater than $\frac{\sum_{(s_0, \dots, s_t^*) \in S_0 \times \dots \times S_t} \prod_{k=0}^t \mu_k^{s_k-1}(s_k) \tau_{s_k}(p_k, A_k)}{\sum_{(s_0, \dots, s_{t-1}) \in S_0 \times \dots \times S_{t-1}} \prod_{k=0}^{t-1} \mu_k^{s_k-1}(s_k) \tau_{s_k}(p_k, A_k)} > 0$, i.e., the conditional probability that s_t^* realizes after history h^{t-1} (see Lemma E.5). Thus, condition (ii) in Axiom C.3 is satisfied.

For the converse, we prove the contrapositive. If (i') does not hold, then by the representation, we have $U_{s_t}(z, A_{t+1}) = U_{s_t}(z, B_{t+1})$ for all $s_t \in S_t^*$. Take menus C'_{t+1}, C_{t+1} such that $U_{s_t}(z, C'_{t+1}) > U_{s_t}(z, C_{t+1})$ for all s_t .⁷⁷ Then define $A_{t+1}^n := \frac{1}{n}C'_{t+1} + \frac{n-1}{n}A_{t+1}$ and $B_{t+1}^n := \frac{1}{n}C'_{t+1} + \frac{n-1}{n}B_{t+1}$ for each n . By construction, $A_{t+1}^n \rightarrow^m A_{t+1}$ and $B_{t+1}^n \rightarrow^m B_{t+1}$. For each s_t , by linearity of $U_{s_t}(z, \cdot)$, it follows that $U_{s_t}(z, A_{t+1}^n) > U_{s_t}(z, B_{t+1}^n)$ for every n . Thus for any $s_t \in S_t^*$, $\frac{1}{2}p_t + \frac{1}{2}(z, B_{t+1}^n) \notin M(\frac{1}{2}A_t + \frac{1}{2}\{(z, A_{t+1}^n), (z, B_{t+1}^n)\}, U_{s_t})$. But then $\rho_t(\frac{1}{2}p_t + \frac{1}{2}(z, B_{t+1}^n); \frac{1}{2}A_t + \frac{1}{2}\{(z, B_{t+1}^n), (z, A_{t+1}^n)\} | h^{t-1}) = 0$ for every n , so that condition (ii) is violated. This completes the proof of necessity.

D Proof of Theorem 3

Instead of proving the two-period characterization of BEB in Theorem 3, this section establishes a generalization of Theorem 3 for arbitrary horizon T . Section D.1 presents the T -period axiom for

⁷⁷Such menus exist by the following argument. Note first that for any $s_{t+1} \in S_{t+1}$, since $U_{s_{t+1}}$ is nonconstant, we can find $g_{t+1}(s_{t+1}), b_{t+1}(s_{t+1}) \in \Delta(X_{t+1})$ such that $U_{s_{t+1}}(g_{t+1}(s_{t+1})) > U_{s_{t+1}}(b_{t+1}(s_{t+1}))$. Let $C'_{t+1} := \{g_{t+1}(s_{t+1}), b_{t+1}(s_{t+1}) : s_{t+1} \in S_{t+1}\}$, and for every s_t , let $A_{t+1}(s_t) := \{b_{t+1}(s_{t+1})\}$ for some $s_{t+1} \in \text{supp} \mu_{t+1}^{s_t}$. Then $U_{s_t}(z, C'_{t+1}) \geq U_{s_t}(z, A_{t+1}(s'_t))$ for all s_t, s'_t , with strict inequality for $s_t = s'_t$. Hence, letting $C_{t+1} := \sum_{s_t \in S_t} \frac{1}{|S_t|} A_{t+1}(s_t)$, linearity implies $U_{s_t}(z, C'_{t+1}) > U_{s_t}(z, C_{t+1})$ for all s_t , as required.

BEB. Sections D.2 and D.3 establish sufficiency and necessity of this axiom.

D.1 Characterization of BEB for Arbitrary T

For any consumption lottery $\ell \in \Delta(Z)$ and menu $A_{t+1} \in \mathcal{A}_{t+1}$, define $(\ell, A_{t+1}) \in \Delta(X_t)$ to be the period- t lottery that yields current consumption according to ℓ and yields continuation menu A_{t+1} for sure; i.e., $(\ell, A_{t+1}) := \sum_{z_t \in Z} \ell(z_t) \delta_{(z_t, A_{t+1})}$. Then for each $t \leq T-1$, $\ell_t, \ell_{t+1} \in \Delta(Z)$, and $A_{t+2} \in \mathcal{A}_{t+2}$, we define $(\ell_t, \ell_{t+1}, A_{t+2}) := (\ell_t, \{p_{t+1}\})$ such that $p_{t+1} = (\ell_{t+1}, A_{t+2})$.⁷⁸

We generalize Axiom 8 and Condition 1 as follows:

Axiom D.1 (Stationary Consumption Preference). For each history h^{t-1} , if $(\ell, \ell, A_{t+2}), (\ell', \ell', A_{t+2}) \in \mathcal{A}_t \in \mathcal{A}_t^*(h^{t-1})$, then

$$\rho_t((\ell, \ell', A_{t+2}); A_t | h^{t-1}) = 0.$$

Condition D.1 (Uniformly Ranked Pair). There exist $\bar{\ell}, \underline{\ell} \in \Delta(Z)$ such that $A_t := \{(\bar{\ell}, A_{t+1}), (\underline{\ell}, A_{t+1})\} \in \mathcal{A}_t^*(h^{t-1})$ and $\rho_t((\bar{\ell}, A_{t+1}); A_t | h^{t-1}) = 1$ for all t, A_{t+1} , and h^{t-1} .

We have the following T -period generalization of Theorem 3:

Theorem D.1. Suppose that ρ admits a BEU representation and Condition D.1 is satisfied. Then ρ satisfies Axioms D.1 if and only if ρ admits a BEB representation.

D.2 Proof of Theorem D.1: Sufficiency

Suppose that ρ admits a BEU representation and that Condition D.1 and Axiom D.1 hold. By Proposition A.1, ρ admits an S-based BEU representation $(S_t, \{\mu_t^{s_t-1}\}_{s_t-1 \in S_{t-1}}, \{U_{s_t}, u_{s_t}, \tau_{s_t}\}_{s_t \in S_t})_{t=0, \dots, T}$. Up to adding appropriate constants to each utility u_{s_t} and U_{s_t} , we can ensure that $\sum_{z \in Z} u_{s_t}(z) = 0$ for all $t = 0, \dots, T$ and $s_t \in S_t$ without affecting that $(S_t, \{\mu_t^{s_t-1}\}_{s_t-1 \in S_{t-1}}, \{U_{s_t}, u_{s_t}, \tau_{s_t}\}_{s_t \in S_t})_{t=0, \dots, T}$ is an S-based BEU representation of ρ . We will show that this representation is in fact an S-based BEB representation, i.e., for each $t \leq T-1$ and s_t , there exists $\delta_{s_t} > 0$ such that we have $u_{s_t} = \frac{1}{\delta_{s_t}} \sum_{s_{t+1}} \mu_{t+1}^{s_t}(s_{t+1}) u_{s_{t+1}}$. By Proposition A.1, this implies that ρ admits a BEB representation.

Condition D.1 ensures that all felicities u_{s_t} agree on the ranking between $\bar{\ell}$ and $\underline{\ell}$:

Lemma D.1. $u_{s_t}(\bar{\ell}) > u_{s_t}(\underline{\ell})$ holds for all t and s_t .

Proof. Consider any t and $s_t \in S_t$ and the state s_{t-1} such that $\mu_t^{s_t-1}(s_t) > 0$. Take a separating history h^{t-1} for s_{t-1} and any A_{t+1} . Let $A_t := \{(\bar{\ell}, A_{t+1}), (\underline{\ell}, A_{t+1})\}$. Then Condition D.1 ensures $A_t \in \mathcal{A}_t^*(h^{t-1})$ and $\rho_t((\bar{\ell}, A_{t+1}); A_t | h^{t-1}) = 1$, which by Lemma E.3 implies $U_{s_t}(\bar{\ell}, A_{t+1}) > U_{s_t}(\underline{\ell}, A_{t+1})$. By the separability of the representation, it follows that $u_{s_t}(\bar{\ell}) > u_{s_t}(\underline{\ell})$. \blacksquare

For any $t = 0, \dots, T-1$ and $s_t \in S_t$, let $u_{s_t}^+ := \sum_{s_{t+1}} \mu_{t+1}^{s_t}(s_{t+1}) u_{s_{t+1}}$ denote the expected period $t+1$ felicity at state s_t . Note that Lemma D.1 ensures that each $u_{s_t}^+$ is non-constant. We show that Axiom D.1 (Stationary Consumption Preference) implies that u_{s_t} and $u_{s_t}^+$ induce the same preference over $\Delta(Z)$:

Lemma D.2. $u_{s_t} \approx u_{s_t}^+$ holds for all $t \leq T-1$ and $s_t \in S_t$.

⁷⁸In the case of $t = T-1$ by abusing notation we are using $(\ell_t, \ell_{t+1}, A_{t+2})$ to denote the lottery that yields ℓ_t in period $T-1$ and ℓ_{t+1} in period T .

Proof. Suppose to the contrary that $u_{s_t^*} \not\approx u_{s_t^*}^+$ for some $s_t^* \in S_t$. Since $u_{s_t^*}$ and $u_{s_t^*}^+$ are nonconstant, there exist $\ell, \ell' \in \Delta(Z)$ such that $u_{s_t^*}(\ell) > u_{s_t^*}(\ell')$ and $u_{s_t^*}^+(\ell) < u_{s_t^*}^+(\ell')$. By slightly perturbing ℓ and ℓ' if needed, we can assume that $u_{s_t}(\ell) \neq u_{s_t}(\ell')$ and $u_{s_t}^+(\ell) \neq u_{s_t}^+(\ell')$ for all $s_t \in S_t$, since all u_{s_t} and $u_{s_t}^+$ are nonconstant.

Fix any A_{t+2} and let $A_t := \{(\ell, \ell, A_{t+2}), (\ell', \ell', A_{t+2}), (\ell, \ell', A_{t+2})\}$. Then, by the separability of the representation, we have that $|M(A_t, U_{s_t})| = 1$ for all $s_t \in S_t$, with unique element given by $(\operatorname{argmax}_{(\ell_t, \ell_{t+1}) \in \{\ell, \ell'\}^2} u_{s_t}(\ell_t) + u_{s_t}^+(\ell_{t+1}), A_{t+2})$. In particular, $M(A_t, U_{s_t^*}) = \{(\ell, \ell', A_{t+2})\}$. Let s_{t-1} be the unique state such that $\mu_t^{s_{t-1}}(s_t^*) > 0$ and take a separating history h^{t-1} for s_{t-1} . Then Lemma E.3 implies that $A_t \in \mathcal{A}_t^*(h^{t-1})$ and $\rho_t((\ell, \ell', A_{t+2}), A_t | h^{t-1}) \geq \mu_t^{s_{t-1}}(s_t^*) > 0$, contradicting Axiom D.1. \blacksquare

Since each u_{s_t} is nonconstant by Lemma D.1, Lemma D.2 implies that for each $t \leq T-1$ and s_t there exist constants $\delta_{s_t} > 0, \gamma_{s_t} \in \mathbb{R}$ such that $u_{s_t}^+ = \delta_{s_t} u_{s_t} + \gamma_{s_t}$. Since we have normalized felicities such that $\sum_{z \in Z} u_{s_t}(z) = 0$ for any t' and $s_{t'}$, we must have $\gamma_{s_t} = 0$. This completes the proof that ρ admits an S-based BEB representation.

D.3 Proof of Theorem D.1: Necessity

Suppose that ρ admits a BEB representation. By Proposition A.1, ρ admits an S-based BEB representation $(S_t, \{\mu_t^{s_{t-1}}\}_{s_{t-1} \in S_{t-1}}, \{U_{s_t}, u_{s_t}, \delta_{s_t}, \tau_{s_t}\}_{s_t \in S_t})_{t=0, \dots, T}$.

To verify Axiom D.1, take any history h^{t-1} and consider $(\ell, \ell, A_{t+2}), (\ell', \ell', A_{t+2}) \in A_t \in \mathcal{A}_t^*(h^{t-1})$. If $\rho_t((\ell, \ell', A_{t+2}), A_t | h^{t-1}) > 0$, then by Lemma E.3, we have $U_{s_t}((\ell, \ell', A_{t+2})) > U_{s_t}((\ell, \ell, A_{t+2}), U_{s_t}((\ell', \ell', A_{t+2})))$ for some s_t . By the representation, this implies that $u_{s_t}(\ell) > u_{s_t}(\ell')$ and $\sum_{s_{t+1}} \mu_t^{s_t}(s_{t+1}) u_{s_{t+1}}(\ell) < \sum_{s_{t+1}} \mu_t^{s_t}(s_{t+1}) u_{s_{t+1}}(\ell')$. But this contradicts the fact that $u_{s_t} = \frac{1}{\delta_{s_t}} \sum_{s_{t+1}} \mu_t^{s_t}(s_{t+1}) u_{s_{t+1}}$ and $\delta_{s_t} > 0$.

E Additional Lemmas

This section collects together several lemmas that are used throughout Sections B–D. The proofs are provided in Supplementary Appendix J.2.

Lemma E.1. For all $t = 0, \dots, T$, X_t is a separable metric space, where $X_T := Z$ is endowed with the discrete metric and for all $t \leq T-1$, we recursively endow $\Delta(X_{t+1})$ with the induced topology of weak convergence, $\mathcal{A}_{t+1} := \mathcal{K}(\Delta(X_{t+1}))$ with the induced Hausdorff topology, and $X_t := Z \times \mathcal{A}_{t+1}$ with the induced product topology.

Lemma E.2. Let Y be any set (possibly infinite) and let $\{U_s : s \in S\} \subseteq \mathbb{R}^Y$ be a collection of nonconstant vNM utility functions indexed by a finite set S such that $U_s \not\approx U_{s'}$ for any distinct $s, s' \in S$. Then there is a collection of lotteries $\{p^s : s \in S\} \subseteq \Delta(Y)$ such that $U_s(p^s) > U_s(p^{s'})$ for any distinct $s, s' \in S$.

Lemma E.3. Fix $t = 0, \dots, T$. Suppose $(S_{t'}, \{\mu_{t'}^{s_{t'-1}}\}_{s_{t'-1} \in S_{t'-1}}, \{U_{s_{t'}}, \tau_{s_{t'}}\}_{s_{t'} \in S_{t'}})$ satisfy DREU1 and DREU2 for all $t' \leq t$. Take any $h^{t-1} \in \mathcal{H}_{t-1}$ and let $S(h^{t-1}) \subseteq S_{t-1}$ denote the set of states consistent with h^{t-1} . Then for any $A_t \in \mathcal{A}_t$, the following are equivalent:

- (i). $A_t \in \mathcal{A}_t^*(h^{t-1})$
- (ii). For each $s_{t-1} \in S(h^{t-1})$ and $s_t \in \operatorname{supp} \mu_t^{s_{t-1}}$, $|M(A_t, U_{s_t})| = 1$.

Lemma E.4. Suppose that ρ satisfies Axiom B.2. Fix $t \geq 1$, $A_t \in \mathcal{A}_t$, $h^{t-1} = (A_0, p_0, \dots, A_{t-1}, p_{t-1}) \in \mathcal{H}_{t-1}$, and $\lambda = (\lambda_n)_{n=0}^{t-1}$, $\hat{\lambda} = (\hat{\lambda}_n)_{n=0}^{t-1} \in (0, 1]^t$. Suppose $d^{t-1} = (\{q_n\}, q_n)_{n=0}^{t-1}$, $\hat{d}^{t-1} = (\{\hat{q}_n\}, \hat{q}_n)_{n=0}^{t-1} \in \mathcal{D}_{t-1}$ satisfy $\lambda h^{t-1} + (1 - \lambda)d^{t-1}$, $\hat{\lambda} h^{t-1} + (1 - \hat{\lambda})\hat{d}^{t-1} \in \mathcal{H}_{t-1}(A_t)$, where $\lambda h^{t-1} + (1 - \lambda)d^{t-1} := (\lambda_n A_n + (1 - \lambda_n)\{q_n\}, \lambda_n p_n + (1 - \lambda_n)q_n)_{n=0}^{t-1}$ and $\hat{\lambda} h^{t-1} + (1 - \hat{\lambda})\hat{d}^{t-1}$ is defined analogously. Then

$$\rho_t(\cdot; A_t | \lambda h^{t-1} + (1 - \lambda)d^{t-1}) = \rho_t(\cdot; A_t | \hat{\lambda} h^{t-1} + (1 - \hat{\lambda})\hat{d}^{t-1}),$$

and hence, $\rho_t^{h^{t-1}}(\cdot; A_t) = \rho_t(\cdot; A_t | \lambda h^{t-1} + (1 - \lambda)d^{t-1})$.

Lemma E.5. Fix $t = 0, \dots, T$. Suppose $(S_{t'}, \{\mu_{s_{t'}}^{s_{t'}-1}\}_{s_{t'}-1 \in S_{t'}-1}, \{U_{s_{t'}}, \tau_{s_{t'}}\}_{s_{t'} \in S_{t'}})$ satisfy DREU1 and DREU2 for all $t' \leq t$. Then the extended version of ρ from Definition 10 also satisfies DREU2 for all $t' \leq t$, i.e., for all $p_{t'}, A_{t'}$, and $h^{t'-1} = (A_0, p_0, \dots, A_{t'-1}, p_{t'-1}) \in \mathcal{H}_{t'-1}$,⁷⁹ we have

$$\rho_{t'}(p_{t'}, A_{t'} | h^{t'-1}) = \frac{\sum_{(s_0, \dots, s_{t'}) \in S_0 \times \dots \times S_{t'}} \prod_{k=0}^{t'} \mu_k^{s_k-1}(s_k) \tau_{s_k}(p_k, A_k)}{\sum_{(s_0, \dots, s_{t'-1}) \in S_0 \times \dots \times S_{t'-1}} \prod_{k=0}^{t'-1} \mu_k^{s_k-1}(s_k) \tau_{s_k}(p_k, A_k)}.$$

Lemma E.6. Fix $t = 0, \dots, T$. Suppose $(S_{t'}, \{\mu_{s_{t'}}^{s_{t'}-1}\}_{s_{t'}-1 \in S_{t'}-1}, \{U_{s_{t'}}, \tau_{s_{t'}}\}_{s_{t'} \in S_{t'}})$ satisfy DREU1 and DREU2 for all $t' \leq t$. Fix any $s_{t-1} \in S_{t-1}$, separating history h^{t-1} for s_{t-1} , and $A_t \in \mathcal{A}_t$. Then there exists a sequence $A_t^n \rightarrow^m A_t$ such that $A_t^n \in \mathcal{A}_{t+1}^*(h^t)$ for all n . Moreover, given any $s_t^* \in \text{supp} \mu_t^{s_t-1}$ and $p_t^* \in M(A_t, U_{s_t^*})$, we can ensure in this construction that there is $p_t^n(s_t^*) \in A_t^n$ with $p_t^n(s_t^*) \rightarrow^m p_t^*$ such that $\mathcal{U}_{s_t}(A_t^n, p_t^n(s_t^*)) = \{U_{s_t^*}\}$ for all n .

⁷⁹For $t' = 0$, we abuse notation by letting $\rho_{t'}(\cdot | h^{t'-1})$ denote $\rho_0(\cdot)$ for all $h^{t'-1}$.

References

- ABDULKADIROGLU, A., J. D. ANGRIST, Y. NARITA, AND P. A. PATHAK (forthcoming): “Research design meets market design: Using centralized assignment for impact evaluation,” *Econometrica*.
- AGUIRREGABIRIA, V., AND P. MIRA (2010): “Dynamic discrete choice structural models: A survey,” *Journal of Econometrics*, 156(1), 38–67.
- AHN, D. S., AND T. SARVER (2013): “Preference for flexibility and random choice,” *Econometrica*, 81(1), 341–361.
- ALIPRANTIS, C. D., AND K. C. BORDER (2006): *Infinite Dimensional Analysis: a Hitchhiker’s Guide*. Springer, Berlin; London.
- ANGRIST, J., P. HULL, P. A. PATHAK, AND C. WALTERS (forthcoming): “Leveraging lotteries for school value-added: Testing and estimation,” *Quarterly Journal of Economics*.
- ANSCOMBE, F. J., AND R. J. AUMANN (1963): “A definition of subjective probability,” *The annals of mathematical statistics*, 34(1), 199–205.
- APESTEGUIA, J., M. BALLESTER, AND J. LU (2017): “Single-Crossing Random Utility Models,” *Econometrica*, 85(2), 661–674.
- APESTEGUIA, J., AND M. A. BALLESTER (2018): “Monotone Stochastic Choice Models: The Case of Risk and Time Preferences,” *Journal of Political Economy*, 126(1), 74–106.
- AUGENBLICK, N., M. NIEDERLE, AND C. SPRENGER (2015): “Working over time: Dynamic inconsistency in real effort tasks,” *The Quarterly Journal of Economics*, p. qjv020.
- BARBERÁ, S., AND P. PATTANAİK (1986): “Falmagne and the rationalizability of stochastic choices in terms of random orderings,” *Econometrica*, pp. 707–715.
- BECKER, G. M., M. H. DEGROOT, AND J. MARSCHAK (1964): “Measuring utility by a single-response sequential method,” *Behavioral science*, 9(3), 226–232.
- BECKER, G. S., AND K. M. MURPHY (1988): “A theory of rational addiction,” *Journal of political Economy*, 96(4), 675–700.
- BLOCK, D., AND J. MARSCHAK (1960): “Random Orderings And Stochastic Theories of Responses,” in *Contributions To Probability And Statistics*, ed. by I. O. et al. Stanford: Stanford University Press.
- CERIGIONI, F. (2017): “Stochastic Choice and Familiarity: Inertia and The Mere Exposure Effect,” working paper.
- CERREIA-VIOGLIO, S., F. MACCHERONI, M. MARINACCI, AND RUSTICHINI (2017): “Multinomial logit processes and preference discovery: inside and outside the black box,” working paper.
- CHAMBERS, C., AND N. LAMBERT (2017): “Dynamically Eliciting Unobservable Information,” .
- DEAN, M., AND J. MCNEILL (2016): “Preference for Flexibility and Random Choice: an Experimental Analysis,” working paper.
- DEKEL, E., B. LIPMAN, AND A. RUSTICHINI (2001): “Representing preferences with a unique subjective state space,” *Econometrica*, 69(4), 891–934.
- DEKEL, E., B. L. LIPMAN, A. RUSTICHINI, AND T. SARVER (2007): “Representing Preferences with a Unique Subjective State Space: A Corrigendum,” *Econometrica*, 75(2), 591–600.
- DEMING, D. J. (2011): “Better schools, less crime?,” *The Quarterly Journal of Economics*, p. qjr036.

- DEMING, D. J., J. S. HASTINGS, T. J. KANE, AND D. O. STAIGER (2014): “School choice, school quality, and postsecondary attainment,” *The American economic review*, 104(3), 991–1013.
- DILLENBERGER, D., J. S. LLERAS, P. SADOWSKI, AND N. TAKEOKA (2014): “A theory of subjective learning,” *Journal of Economic Theory*, 153, 287–312.
- DUBÉ, J.-P., G. J. HITSCH, AND P. E. ROSSI (2010): “State dependence and alternative explanations for consumer inertia,” *The RAND Journal of Economics*, 41(3), 417–445.
- DURAJ, J. (2018): “Dynamic Random Subjective Expected Utility,” working paper.
- EPSTEIN, L. G. (1999): “A definition of uncertainty aversion,” *The Review of Economic Studies*, 66(3), 579–608.
- EPSTEIN, L. G., AND S. ZIN (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57(4), 937–969.
- FALMAGNE, J. (1978): “A representation theorem for finite random scale systems,” *Journal of Mathematical Psychology*, 18(1), 52–72.
- FISHBURN, P. (1970): *Utility theory for decision making*. Wiley.
- FISHBURN, P. C. (1984): “On Harsanyi’s Utilitarian Cardinal Welfare Theorem,” *Theory and Decision*, 17, 21–28.
- FRICK, M., R. IJIMA, AND T. STRZALECKI (2017): “Dynamic Random Utility,” Cowles Foundation Discussion Paper # 2092.
- FUDENBERG, D., P. STRACK, AND T. STRZALECKI (2016): “Speed, Accuracy, and the Optimal Timing of Choices,” .
- FUDENBERG, D., AND T. STRZALECKI (2015): “Dynamic logit with choice aversion,” *Econometrica*, 83(2), 651–691.
- GILBOA, I. (1987): “Expected utility with purely subjective non-additive probabilities,” *Journal of mathematical Economics*, 16(1), 65–88.
- GITTINS, J. C., AND D. M. JONES (1972): *A dynamic allocation index for the sequential design of experiments*. University of Cambridge, Department of Engineering.
- GOWRISANKARAN, G., AND M. RYSMAN (2012): “Dynamics of consumer demand for new durable goods,” *Journal of political Economy*, 120(6), 1173–1219.
- GUL, F., P. NATENZON, AND W. PESENDORFER (2014): “Random Choice as Behavioral Optimization,” *Econometrica*, 82(5), 1873–1912.
- GUL, F., AND W. PESENDORFER (2001): “Temptation and self-control,” *Econometrica*, 69(6), 1403–1435.
- GUL, F., AND W. PESENDORFER (2004): “Self-control and the theory of consumption,” *Econometrica*, 72(1), 119–158.
- GUL, F., AND W. PESENDORFER (2006): “Random expected utility,” *Econometrica*, 74(1), 121–146.
- HAUSMAN, J., AND D. MCFADDEN (1984): “Specification tests for the multinomial logit model,” *Econometrica*, pp. 1219–1240.
- HECKMAN, J. J. (1981): “Heterogeneity and state dependence,” in *Studies in labor markets*, pp. 91–140. University of Chicago Press.
- HENDEL, I., AND A. NEVO (2006): “Measuring the implications of sales and consumer inventory

- behavior,” *Econometrica*, 74(6), 1637–1673.
- HOTZ, V. J., AND R. A. MILLER (1993): “Conditional choice probabilities and the estimation of dynamic models,” *The Review of Economic Studies*, 60(3), 497–529.
- HU, Y., AND M. SHUM (2012): “Nonparametric identification of dynamic models with unobserved state variables,” *Journal of Econometrics*, 171(1), 32–44.
- HUBER, J., J. W. PAYNE, AND C. PUTO (1982): “Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis,” *Journal of Consumer Research*, 9.
- JEULAND, A. P. (1979): “Brand choice inertia as one aspect of the notion of brand loyalty,” *Management Science*, 25(7), 671–682.
- KASAHARA, H., AND K. SHIMOTSU (2009): “Nonparametric identification of finite mixture models of dynamic discrete choices,” *Econometrica*, 77(1), 135–175.
- KE, S. (2018): “Rational expectation of mistakes and a measure of error-proneness,” *Theoretical Economics*, 13(2), 527–552.
- KEANE, M. P. (1997): “Modeling heterogeneity and state dependence in consumer choice behavior,” *Journal of Business & Economic Statistics*, 15(3), 310–327.
- KENNAN, J., AND J. WALKER (2011): “The effect of expected income on individual migration decisions,” *Econometrica*, 79(1), 211–251.
- KHAMSI, M. A., AND W. A. KIRK (2011): *An introduction to metric spaces and fixed point theory*, vol. 53. John Wiley & Sons.
- KITAMURA, Y., AND J. STOYE (2018): “Nonparametric analysis of random utility models,” *Econometrica* forthcoming.
- KREPS, D. (1979): “A representation theorem for” preference for flexibility,” *Econometrica*, pp. 565–577.
- KREPS, D., AND E. PORTEUS (1978): “Temporal Resolution of Uncertainty and Dynamic Choice Theory,” *Econometrica*, 46(1), 185–200.
- KRISHNA, R. V., AND P. SADOWSKI (2014): “Dynamic preference for flexibility,” *Econometrica*, 82(2), 655–703.
- KRISHNA, V., AND P. SADOWSKI (2016): “Randomly Evolving Tastes and Delayed Commitment,” *mimeo*.
- LU, J. (2016a): “A Bayesian Theory of State-Dependent Utilities,” working paper.
- (2016b): “Random choice and private information,” *Econometrica*, 84(6), 1983–2027.
- LU, J., AND K. SAITO (2018a): “Random intertemporal choice,” *Journal of Economic Theory*, 177, 780–815.
- (2018b): “Repeated Choice,” *mimeo*.
- LUCE, D. (1959): *Individual choice behavior*. John Wiley.
- MA, W. (2018): “Random expected utility theory with a continuum of prizes,” *Annals of Operations Research*, pp. 1–23.
- MAGNAC, T., AND D. THESMAR (2002): “Identifying dynamic discrete decision processes,” *Econometrica*, 70(2), 801–816.
- MANZINI, P., AND M. MARIOTTI (2018): “Dual random utility maximisation,” *Journal of Economic*

Theory.

- MARTINEZ, S.-K., S. MEIER, AND C. SPRENGER (2017): “Procrastination in the Field: Evidence from Tax Filing,” .
- McFADDEN, D., AND M. K. RICHTER (1990): “Stochastic rationality and revealed stochastic preference,” in *Preferences, Uncertainty, and Optimality: Essays in Honor of Leo Hurwicz*, pp. 161–186. Westview Press: Boulder, CO.
- MILLER, R. A. (1984): “Job matching and occupational choice,” *The Journal of Political Economy*, pp. 1086–1120.
- NOOR, J. (2011): “Temptation and Revealed Preference,” *Econometrica*, 79(2), 601–644.
- NORETS, A., AND X. TANG (2013): “Semiparametric Inference in dynamic binary choice models,” *The Review of Economic Studies*, p. rdt050.
- PAKES, A. (1986): “Patents as options: Some estimates of the value of holding European patent stocks,” *Econometrica*, 54, 755–784.
- PENNESI, D. (2017): “Intertemporal discrete choice,” working paper.
- RAO, K. P. S. B., AND M. B. RAO (2012): *Theory of Charges: A Study of Finitely Additive Measures*. Academic Press.
- ROBBINS, H. (1952): “Some aspects of the sequential design of experiments,” *Bulletin of the American Mathematical Society*, 58(5), 527–535.
- RUST, J. (1989): “A Dynamic Programming Model of Retirement Behavior,” in *The Economics of Aging*, ed. by D. Wise, pp. 359–398. University of Chicago Press: Chicago.
- (1994): “Structural estimation of Markov decision processes,” *Handbook of econometrics*, 4, 3081–3143.
- SAVAGE, L. J. (1972): *The foundations of statistics*. Courier Corporation.
- SEETHARAMAN, P. (2004): “Modeling multiple sources of state dependence in random utility models: A distributed lag approach,” *Marketing Science*, 23(2), 263–271.
- SHAKED, M., AND J. G. SHANTHIKUMAR (2007): *Stochastic orders*. Springer Science & Business Media.
- SWEETING, A. (2013): “Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry,” *Econometrica*, 81(5), 1763–1803.
- TOUSSAERT, S. (2018): “Eliciting Temptation and Self-Control Through Menu Choices: A Lab Experiment,” *Econometrica*, 86(3), 859–889.
- WHITE, H. (1982): “Maximum likelihood estimation of misspecified models,” *Econometrica*, pp. 1–25.
- WILCOX, N. T. (2011): “‘Stochastically more risk averse:’ A contextual theory of stochastic discrete choice under risk,” *Journal of Econometrics*, 162(1), 89–104.
- ZAJONC, R. B. (2001): “Mere exposure: A gateway to the subliminal,” *Current directions in psychological science*, 10(6), 224–228.