RANDOM UTILITY, REPEATED CHOICE, AND CONSUMPTION DEPENDENCE

CHRISTOPHER TURANSICK

Abstract. We study the repeated choice interpretation of stochastic choice and the individual interpretation of the random utility model. We consider a myopic agent whose distribution over preferences tomorrow potentially depends on their consumption and preference today. Even when the agent is classically rational in each of their static decisions, there are forms of consumption dependence which are inconsistent with random utility as a model of intertemporal aggregation. We offer two characterizations of Markovian consumption dependence which are consistent with random utility. Further, we characterize the behavioral content of consumption dependence in a two period stochastic choice model.

1. Introduction

The stochastic choice paradigm is one often considered in modern choice theory as it allows researchers to model data often seen in empirical applications. Stochastic choice data is frequently assumed to arise from some underlying distribution over preferences which governs choice. This is called the random utility model and it, along with stochastic choice, has two common interpretations. The first interpretation is that stochastic choice is the aggregate choice of a market of heterogeneous agents. In this case, the random utility model can be thought of as modeling a distribution over heterogeneous agents, each of whom is classically rational. Borrowing language from

Date: May 30, 2023.
I am grateful to Sean Horan, Kyle Monk, and John Rehbeck for their helpful comments during the course of this project. I am especially grateful to Peter Caradonna, Christopher Chambers, and Yusufcan Masatlioglu for their continued support and insightful conversations throughout the course of this project.
Turansick: Department of Economics, Georgetown University, ICC 580 37th and O Streets NW, Washington DC 20057. E-mail: `cmt152@georgetown.edu` .

Lu and Saito (2020), we call this interpretation the *population interpretation*. The other interpretation of stochastic choice is that it is the result of aggregating the repeated decisions of a singular decision maker across time. Here the random utility model is thought of as modeling the stationary distribution of some stochastic process over classically rational preferences. We call this interpretation the *individual interpretation*. We study the individual interpretation of the random utility model in order to see if it is valid when the underlying stochastic process between preferences is subject to consumption dependence.

There are at least two difficulties present in intertemporal aggregation that are not present when aggregating across agents. When making repeated decisions, an agent may be dynamically sophisticated and take future periods into account. If this is the case, an agent's choice does not represent maximization of a preference over today's alternatives but rather maximization of a Bellman equation that takes into account today's utility as well as the agent's continuation value. Lu and Saito (2020) study an environment where a potentially dynamically sophisticated agent makes repeated decisions and provides examples of where this dynamic sophistication can cause issues when performing inference. A second difficulty that may arise, even in the case of a myopic agent, is that the agent's preference tomorrow may be a random function of not only their preference today but of their choice today as well. Borrowing the language of Frick et al. (2019), this phenomenon is called *consumption dependence*. In the case of consumption dependence, the agent's stationary distribution over preferences may depend on the set of alternatives available to them. We study this case and our primary goal is to characterize which forms of (Markovian) consumption dependence lead to the agent's stationary distribution over preferences being independent of the set of alternatives they face.

There are a multitude of environments and causes which can induce consumption dependence. One such cause is habit formation or addiction as discussed in Becker and Murphy (1988) and Gul and Pesendorfer (2007). With habit formation, the consumption of a good today increases the relative utility of that good tomorrow. Thus, when considering habit formation, the choice of a good today would lead to a high probability of drawing a preference in favor of that good tomorrow. Due to this, when looking at the repeated choice of a single agent, we would expect to see long periods of time where the same choice is made. Alternatively, when facing repeated decisions, an agent may have a preference for variation. In this case, consuming a good today would lead to the relative utility of that good decreasing tomorrow. In this setting we would expect the repeated decisions of an agent to exhibit low frequency of sequential choices being the same. Preference for variation can be thought of as a dynamic foundation for preferences for randomization as discussed in Machina (1985), Agranov and Ortoleva (2017), and Cerreia-Vioglio et al. (2019).

We model consumption dependence through the use of a transition function t. Given today's choice x and preference $\succ$, $t(x, \succ)$ returns a distribution over preferences. The next time the agent faces a decision problem, they draw a preference according to $t(x, \succ)$ and maximize their realized preference. We call this model the Markov random utility model. The Markov random utility model allows for habit formation, preference for variation, and any arbitrary form of single period consumption dependence.[1] For each choice set, the transition function t induces a unique Markov chain over preferences. When the state is $\succ$, the transition probabilities of this Markov chain are defined by $t(x, \succ)$ where x is the maximal element of $\succ$ in the agent's choice set. This Markov chain may vary between choice sets as the choice of x can never occur in a choice set that does not include x. Our primary goal is to characterize which transition functions induce the same stationary distribution over preferences at each choice set. We call such transition functions *stable*. We are interested in stable transition functions as this is the case when the Markov random utility model is consistent with the individual interpretation of the random utility model.

Our main results, Theorem 1 and Theorem 2, characterize exactly which transition functions are stable. Theorem 1 does so through the use of a martingale style condition. Theorem 1 says that a transition function is stable if and only if for every pair of choice

[1]While the Markov random utility model restricts itself to dependence only on the prior period, the techniques we use in this paper can be extended to handle dependence on the prior n periods for any finite and fixed n.

sets A and $A \setminus \{x\}$, the transition probabilities from preferences which choose x in A are on average the same in A and in $A \setminus \{x\}$. Theorem 2 characterizes stability through a no money pump condition. We call a zero value bet a bet that costs as much as it pays conditional on the realization of the agent's preference. The interpretation is that some observer is betting on the agent's preference tomorrow given the agent's preference today and is potentially doing so at multiple choice sets. No money pump asks that there is no zero value bet that gives a strict expected profit for every realization of the agent's preference today. Our Theorem 2 says that a transition function is stable if and only if it satisfies no money pump.

In Theorem 1 and Theorem 2, we make the assumption that the agent repeatedly makes choices from the same choice set. In Theorem 3, we drop this assumption and allow the agent's choice set to vary over time according to some ergodic Markov chain. Theorem 3 characterizes stability through an extension of the no money pump condition of Theorem 2. The main difference between the two conditions is that in Theorem 2, the expected revenue of the bet is calculated assuming that probability of a preference in choice set A tomorrow is inherited directly from choice set A today. However, in Theorem 3 the probability of a preference in choice set A tomorrow is inherited from every choice set according to the stationary distribution of the Markov chain varying the agent's choice set.

In Theorem 4 and Theorem 5, we consider a two period stochastic choice model. In the first period choices are made according to the random utility model. Given the realized preference and choice in the first period, preferences then transition according to some transition function t in the second period. Choices in the second period capture the dynamic properties of the Markov random utility model. Theorem 4 characterizes this model when the transition function only depends on the agent's choice. Theorem 5 characterizes the general version of this model.

The rest of this paper is organized is follows. We conclude this section with a review of the related literature. In Section 2, we formally introduce our model. In Section 3, we introduce and discuss our main results, Theorem 1 and Theorem 2. In Section

4, we generalize our baseline model to allow choice sets to vary over time and present Theorem 3. In Section 5, we introduce and characterize a two period stochastic choice model which captures the dynamic properties of the Markov random utility model. We conclude in Section 6.

1.1. **Related Literature.** Our paper is most closely related to Lu and Saito (2020) and Frick et al. (2019). Lu and Saito (2020) studies the individual interpretation of stochastic choice when the agent's utility potentially takes into account both today's consumption as well as the agent's continuation value. They show that if the analyst fails to take this dynamic sophistication into account, then estimation of key parameters such as risk aversion or costs in a dynamic discrete choice setting may be biased. In their setting, they find that dynamic sophistication leads to biased estimation if and only if the agent's utility function over today's consumption c and their continuation value v does not take the form $U(c, v) = (1-\beta)u(c)+\beta v$ where u is a random expected utility function and $\beta \in (0, 1]$ is a random discount factor.

Frick et al. (2019) proposes the dynamic random expected utility model in order to axiomatically study choice data that is apparently history dependent. One of their main goals is to characterize patterns of history dependence which are spurious. In the two period formulation of their main model, they consider a joint distribution over random expected utility functions. The first utility function dictates choice in the first period and the second dictates choice in the second period. Despite the lack of consumption dependence, this model still allows for some forms of history dependence. In an earlier version of their paper, they extend their model to explicitly allow for consumption dependence and characterize which forms of history dependence can arise due to consumption dependence. Our paper is concerned with how consumption dependence, as studied in Frick et al. (2019), may lead to troubles when aggregating the repeated decisions of a single agent, a similar problem to the one studied in Lu and Saito (2020).

Closely related to the model of our paper are Li (2022), Chambers et al. (2022), and Kashaev and Aguiar (2022). Each of these papers study the special case of our model in which the transition function of an agent does not depend on their choice and only depends on their preference. We call this case pure preference dependence. All three papers show that the set of inequalities we use to characterize our model is necessary in their environment. Li (2022) also offers a full characterization of pure preference dependence in the style of Clark (1996). While Chambers et al. (2022) does study pure preference dependence, their main focus is on joint stochastic choice with well-defined marginal choice probabilities. Finally, Kashaev and Aguiar (2022) studies the case of pure preference dependence in the price-wealth domain rather than in the abstract domain. In doing so they are able to connect the methodologies of McFadden and Richter (1990) with those of Afriat (1967). This allows the authors to study data that is both heterogeneous across time as well as across agents. Further, Kashaev and Aguiar (2022) study a more general type of data than we do. We assume that choice is observed on all non-empty subsets of a finite choice set. In the price wealth domain of Kashaev and Aguiar (2022) there are a continuum of possible choices and it is only assumed that choice is observed on some set of budgets rather than on every budget.

More generally our paper is related to three strands of literature: (i) random utility and dynamic random utility, (ii) consumption dependence, and (iii) common priors and feasible distributions over posteriors. There is a large literature on the static random utility model of Block and Marschak (1959). Falmagne (1978), Barberá and Pattanaik (1986), McFadden and Richter (1990), and Fiorini (2004) offer axiomatic characterization of the random utility model. Gul and Pesendorfer (2006) introduces the random expected utility model which serves as the basis for many of the papers, including Frick et al. (2019) and Lu and Saito (2020), studying random utility in dynamic settings. To our knowledge, Fudenberg and Strzalecki (2015) is the first to study the axioms of a random utility model in a dynamic setting. Duraj (2018) extends many of the results of Frick et al. (2019) to a setting with a subjective beliefs about the state of the world. Strack and Taubinsky (2021) studies when an expected utility preference in one period can be induced by aggregating a random utility in a second period. Deb and Renou (2021) and Alós-Ferrer and Mihm (2021) consider dynamic environments where agents can learn from past choices. In Alós-Ferrer and Mihm

(2021) agents only learn from their past choices while in Deb and Renou (2021) agents are able to learn from the choices of others.

Our paper is also related to the strand of literature studying consumption dependence. The idea to distinguish between history dependence and consumption dependence dates back to at least Heckman (1978) and Heckman (1981). As mentioned prior, consumption dependence can arise due to many different factors. Becker and Murphy (1988) and Gul and Pesendorfer (2007) study consumption dependence that arises due to addiction. There is also a long literature on consumption dependence that arises due to consumers learning about product characteristics. To name a few, Erdem and Keane (1996), Crawford and Shum (2005), and Pakes et al. (2021) each study environments where a consumer gains knowledge about a product by consuming said product. Chen (2008) and Chen and Risen (2010) study how cognitive dissonance can lead to consumption dependence and the difficulties distinguishing choices that arise from this cognitive dissonance and choices that arise due to some fixed underlying preference.

Lastly, our paper and concept of stability are related to the common prior assumption and the study of feasible distributions of posteriors. Stability asks that the stationary distribution over preferences is the same at each choice set. The common prior assumption asks that each agent in a group have the same prior beliefs about the state of the world. Our Theorem 1 characterizes stability through a martingale condition. In the single agent case, a distribution over posteriors is consistent with a prior if and only if the prior and distribution over posteriors satisfy an analogous martingale condition. In the case of multiple agents, the existence of a common prior has been characterized via no-trade conditions. Milgrom and Stokey (1982) first notes that, in the case of a common prior, agents will not engage in speculative trade. Morris (1994) proves a converse of this statement and shows that agents have a common prior if and only if they will not engage in speculative trade. This characterization is also studied in Feinberg (2000) and Samet (1998). Recently, Arieli et al. (2021) studies the question of feasible joint posterior beliefs.[2] In this setting, the common prior assumption is still

[2]Arieli et al. (2021) restricts to the case of two underlying states. Lang (2022) considers the problem for an any finite state space and is able to show an equivalence between the feasibility of posteriors

characterized by a type of no-trade theorem. Morris (2020) shows how to derive the condition developed in Arieli et al. (2021) from the result of Morris (1994). In doing so, Morris (2020) shows that there is an equivalence between the no-trade condition and a no money pump condition. Our Theorem 2 and Theorem 3 characterize stability through the use of a no money pump condition.

2. Model

2.1. **Preliminaries.** Let X be a finite set of alternatives with typical elements x, y, and z. $\mathcal{L}(X)$ denotes the set of linear orders of X with typical element $\succ$. We let $\Delta(\mathcal{L}(X))$ denote the set of probability distributions over $\mathcal{L}(X)$ with typical element ν. Further, we let $int\Delta(\mathcal{L}(X))$ denote the set of full support probability distributions over $\mathcal{L}(X)$. Define $N(x, A) = \{\succ \in \mathcal{L}(X) | x \succ A \setminus \{x\}\}$ and $M(x, A) = \{\succ \in \mathcal{L}(X) | X \setminus A \succ x \succ A \setminus \{x\}\}$. $N(x, A)$ is the set of linear orders which are maximized by x in A and $M(x, A)$ is the set of linear orders which are maximized by x in A but are not maximized by x in any strict superset of A. Let $\mathcal{X}$ denote the collection of subsets of X with at least two elements with typical elements A and B. Let $x(A, \succ)$ denote the alternative $x \in A$ such that $x \succ y$ for all $y \in A \setminus \{x\}$.

Definition. We call a function $t : X \times \mathcal{L}(X) \to \Delta(\mathcal{L}(X))$ a *transition function*. Further we call a function $t : X \times \mathcal{L}(X) \to int\Delta(\mathcal{L}(X))$ a *strict transition function*.

We use $t_{\succ'}(x, \succ)$ to denote the probability of $\succ'$ given that $(x, \succ)$ is the input to the transition function.

Definition. A function $p : X \times \mathcal{X} \to [0, 1]$ is a random choice rule (rcr) if it satisfies the following.

- $x \notin A \implies p(x, A) = 0$
- $\sum_{x \in A} p(x, A) = 1$

and the existence of a solution to a given multi-commodity flow problem. Lang (2022) also shows that this multi-commodity flow problem reduces to a single commodity flow problem when the state space is binary.

2.2. **Stability and Markov Random Utility.** Our first goal in this paper is to explore the relationship between the random utility model and repeated choice where preferences evolve according to a transition function t. A strict transition function t induces an ergodic Markov chain for each choice set A. We are interested in the stationary distributions of these Markov chains.

For each choice set A, a preference $\succ$ induces the choice $x(A, \succ)$. Consider a Markov matrix M_A whose rows and columns are indexed by the elements of $\mathcal{L}(X)$. Let $\succ$ index the rows and $\succ'$ index the columns. M_A is then given by $\left[m_A(\succ, \succ') = t_{\succ'}(x(A, \succ), \succ)\right]$ where $m_A(\succ, \succ')$ indicates the element of M_A indexed by $(\succ, \succ')$. When t is a strict transition function, M_A represents an ergodic Markov chain with a unique stationary distribution ν_A. Each ν_A then defines choice in each subset. Specifically, we have the following.

$$p(x, A) = \sum_{\succ \in N(x,A)} \nu_A(\succ)$$

We call this model the Markov random utility model (MRUM). We interpret MRUM as follows. There is a myopic agent who repeatedly makes choice decisions from a choice set A. Each time the agent makes a choice, they choose the alternative in A which is ranked highest by their current preference. The agent's preference then updates according to $t(x(A, \succ), \succ)$.

Recall that the standard random utility model (RUM) supposes that there exists some $\nu \in \Delta(\mathcal{L}(X))$ such that choices follow the following rule.

$$p(x, A) = \sum_{\succ \in N(x,A)} \nu(\succ)$$

MRUM and RUM differ in the fact that in RUM ν is the same for each choice set while in MRUM ν_A is allowed to vary between choice sets according to the underlying transitions function. In this sense, RUM is a special case of MRUM.

Definition. A strict transition function t is *stable* if for each $A, B \in \mathcal{X}$, we have $\nu_A = \nu_B$.

Stability exactly captures the cases when MRUM is equivalent to RUM. The individual interpretation of RUM is that stochastic choice captures the repeated choice of a single agent and the underlying distribution ν is the stationary distribution over preferences for some underlying stochastic process. We argue that this underlying stochastic process may depend on the choices of the agent and our goal is to see when this consumption dependence is consistent with RUM.

3. Stability and Stationary Distributions

3.1. **Two Extreme Cases.** Before studying the problem of stability for general strict transition functions, we consider two special cases: pure preference dependence and pure consumption dependence.

Definition. We say that a transition function is *purely preference dependent* if, for all $x, y \in X$ and for all $\succ \in \mathcal{L}(X)$, we have $t(x, \succ) = t(y, \succ)$. In the case of pure preference dependence we write $t(\succ)$ instead of $t(x, \succ)$.

Pure preference dependence captures how analysts have classically thought of the individual interpretation of RUM.[3] Pure preference dependence simply asks that an agent's preference tomorrow only depends on their preference today. This case is consistent with the latent variable hypothesis which assumes that an agent has a fixed utility function but that utility function takes as input the realization of some random variable which is unobserved to the analyst. As an example, an agent's decision to carry an umbrella is determined by the weather, but the weather may not be a part of an analyst's data set.

Observation 1. *Suppose a strict transition function t is purely preference dependent. t is stable.*

Observation 1 tells us that every purely preference dependent transition function is consistent with RUM. This means that the individual interpretation of RUM is valid

[3] In Lu and Saito (2020), an agent's utility is determined by an ergodic utility process. This utility process is a Markov process between expected utility functions.

when one wants to test the latent variable hypothesis. This result follows from the fact that a purely preference dependent transition function takes no input from the choice set of the agent.

Definition. We say that a transition function is *purely consumption dependent* if, for all $x \in X$ and for all $\succ, \succ' \in \mathcal{L}(X)$, we have $t(x, \succ) = t(x, \succ')$. In the case of pure consumption dependence we write $t(x)$ instead of $t(x, \succ)$.

Pure consumption dependence is the opposite extreme of pure preference dependence. In the case of pure consumption dependence, the agent's preference tomorrow is determined only by their choice today. While pure consumption dependence does not allow for the latent variable hypothesis, it allows for other interesting dynamic behaviors. One such case is habit formation, a case which is not consistent with pure preference dependence.

Proposition 1. *Suppose that $|X| \geq 3$ and that a strict transition function t is purely consumption dependent. t is stable if and only if $t(x) = t(y)$ for all $x, y \in X$.*

We leave all proofs to the appendix. Proposition 1 is a negative result. It tells us that pure consumption dependence is consistent with RUM if and only if there is no consumption dependence. In other words, the condition in Proposition 1 asks that the distribution over an agent's preferences tomorrow is independent of their choice today. The main take away from Proposition 1 is that the individual interpretation of RUM may not be valid when we have consumption dependence. Pure preference dependence and pure consumption dependence are the two extreme cases of MRUM and coincide with two extreme relationships with stability; anything goes and impossibility. For general transition functions, there will be some forms of consumption dependence that are still consistent with RUM.

3.2. **A Martingale Condition.** We just saw that, in the case of pure consumption dependence, stability holds if and only if the transition function is the same for each alternative x and y. In the general case, this need not be the case. However, it will be the case that the transition functions must on average remain the same as we move between choice sets. To be formal, as we go from the set A to $A \setminus \{x\}$, the preferences which chose x in A are now choosing some other alternative $y \in A \setminus \{x\}$. Thus, in the case when $t(x, \succ)$ governed transitions in A, $t(y, \succ)$ now governs transitions in $A \setminus \{x\}$. The following condition, which we call *local stability*, captures the idea of transition functions being the same on average.

Definition. Let ν_A be the stationary distribution for some set $A \in \mathcal{X}$. A transition function t is *locally stable at* A if, for all $x \in B \in \mathcal{X}$ with $|B| \geq 3$, we have the following.

$$\sum_{\succ \in N(x,A)} \nu_A(\succ) t(x, \succ) = \sum_{y \in B \setminus \{x\}} \sum_{\succ \in N(x,A) \cap N(y, A \setminus \{x\})} \nu_A(\succ) t(y, \succ) \tag{1}$$

The left hand side of Equation (1) captures the probability the agent chooses x from A today times the distribution over the agent's preference tomorrow when they choose x from choice set A today. The right hand side of Equation (1) captures the probability that the agent has drawn a preference that is maximized by x in A tomorrow but has to choose an alternative from $A \setminus \{x\}$ times the distribution over the agent's preference tomorrow conditional on the prior described event.

Theorem 1. *Given a strict transition function t, the following are equivalent.*

(1) *t is stable.*
(2) *t is locally stable for all $A \in \mathcal{X}$.*
(3) *t is locally stable for some $A \in \mathcal{X}$.*

Theorem 1 tells us that stability and local stability are equivalent. The appeal of local stability is that it tells us, when we compare the stationary distributions of two choice sets, exactly which parts of the stationary distributions matter for stability. We simply need to check the preferences which induce different choices in A and $A \setminus \{x\}$.

3.3. **A No Money Pump Condition.** Theorem 1 provides a simple characterization of stability. However, the characterization requires first calculating the stationary distribution of some choice set in $\mathcal{X}$. We now characterize stability through a no money pump condition which does not require the calculation of any stationary distribution.

Consider an outside observer who wishes to make a bet about the agent's preference tomorrow given the agent's preference today. This outside observer can make this bet at each choice set. We restrict attention to what we call zero value bets.

Definition. We call a function $b : \mathcal{X} \times \mathcal{L}(X) \to \mathbb{R}^+$ a *zero value bet.*

The outside observer is restricted to bets that cost $b(A, \succ)$ to make at choice set A when the agent's preference is $\succ$ but also pay $b(A, \succ)$ at choice set A when the agent's preference is $\succ$. If we restrict ourselves to a single choice set, the outside observer makes a zero expected profit as the bet is repeated. When we look at multiple choice sets concurrently, this is also the case if the outside observer places the bet in every period. If, instead, we allow the outside observer to make the bet only when each choice set agrees on the agent's preference, the outside observer has the opportunity to make a strict profit. In this case, we use $b(\succ) = \sum_{A \in \mathcal{X}} b(A, \succ)$ to denote the total cost of placing bet b when the agent's preference is $\succ$.

Definition. We say that a transition function t satisfies *no money pump* if for each zero value bet b the following holds.

$$\min_{\succ \in \mathcal{L}(X)} \sum_{(A, \succ') \in \mathcal{X} \times \mathcal{L}(X)} b(A, \succ') t_{\succ'}(x(A, \succ), \succ) - b(\succ) \leq 0 \tag{2}$$

The expression $b(\succ)$ in Equation (2) is the total cost of placing bet b when the agent's preference is $\succ$. In the expression on the left in Equation (2), $t_{\succ'}(x(A, \succ), \succ)$ denotes the probability that tomorrow's preference is $\succ'$ at choice set A. This means that the term on the left in Equation (2) captures the expected revenue of bet b tomorrow given that the agent's preference is $\succ$ today. No money pump asks that the outside observer is unable to make a strict profit for every preference $\succ$. Our next theorem shows that no money pump is a tight characterization of stability.

Theorem 2. *A strict transition function t is stable if and only if it satisfies no money pump.*

4. Choice Set Arrival

Thus far we have assumed that an agent's choice set is fixed. However, this differs from reality where an agent's choice set may vary over time and necessarily does so if an analyst has data for multiple choice sets. In this section we extend MRUM to allow the agent's choice set to vary according to an ergodic Markov chain. Our goal is to characterize transition function and choice set arrival process pairs which are stable.

Definition. A function $s : \mathcal{X} \to \Delta(\mathcal{X})$ is an *arrival function* if s represents the transition probabilities of an ergodic Markov chain.

We use $s_B(A)$ to denote the transition probability from A to B. Taken together, a transition function t and an arrival function s jointly define a Markov chain over preference and choice set pairs, $(\succ, A)$. We say that this Markov chain is defined by st as the transition probability from $(\succ, A)$ to $(\succ', B)$ is given by $s_B(A)t_{\succ'}(x(\succ, A), \succ)$. Let π be a typical element of $\Delta(\mathcal{X})$. Further, we let ψ to denote a typical element of $\Delta(\mathcal{X} \times \mathcal{L}(X))$, the set of probability distribution over preference and choice set pairs.

Definition. We say that a strict transition function and arrival function pair (t, s) are *jointly stable* if the unique stationary distribution ψ of the Markov chain defined by st can be written as $\psi(\succ, A) = \nu(\succ)\pi(A)$ for some $\nu \in \Delta(\mathcal{L}(X))$ and $\pi \in \Delta(\mathcal{X})$.

Joint stability simply asks that the stationary distribution over preferences does not depend on the choice set. Of use to us will be the probability of yesterday's choice set given today's choice set in the stationary distribution of s. We use the notation $\rho(B_{\tau-1}|A_\tau)$ to denote the probability in the stationary distribution that last period's choice set was B given that this period's choice set is A. Formally, given the stationary distribution π over choice sets, $\rho(B_{\tau-1}|A_\tau)$ is given by the following.

$$\rho(B_{\tau-1}|A_\tau) = \frac{\pi(B)s_A(B)}{\sum_{C \in \mathcal{X}} \pi(C)s_A(C)} = \frac{\pi(B)s_A(B)}{\pi(A)}$$

Consider the following condition.

Definition. We say that a transition function and arrival function pair (t, s) satisfies *no money pump* if for each zero value bet b the following holds.

$$\min_{\succ \in \mathcal{L}(X)} \sum_{(A,\succ') \in \mathcal{X} \times \mathcal{L}(X)} \sum_{B \in \mathcal{X}} b(A, \succ') \rho(B_{\tau-1}|A_\tau) t_{\succ'}(x(B, \succ), \succ) - b(\succ) \leq 0 \tag{3}$$

Equation (3) has a similar interpretation to that of Equation (2). The expression on the right captures the cost of placing a zero value bet when the preference at each choice set is $\succ$. The expression on the left captures the expected revenue of the zero value bet when the preference at each choice set is $\succ$. The major difference between Equation (2) and Equation (3) is that in Equation (2) the preference of choice set A tomorrow is inherited only from A. However, in Equation (3), the preference of choice set A tomorrow is inherited from every choice set according to the stationary distribution of the arrival function s. As before, Equation (3) asks that no zero value bet makes a strict profit for every preference $\succ$.

Theorem 3. *A strict transition function and arrival function pair (t, s) is jointly stable if and only if the pair satisfies no money pump.*

Similar to Theorem 2, Theorem 3 tells us that stability is equivalent to no money pump. Thus far we have provided tests for stability that rely on observation of a transition function. In empirical settings, choices are observed while preferences are unobserved. This further means that the probability that the agent's preference is $\succ$ tomorrow given that their preference is $\succ'$ today is unobserved. In addition to choices, there are environments where the agent's choice set is observable. This leads us to ask if we can guarantee stability when we only observe a choice set arrival function s.

Definition. We say that an arrival function s is *stable* if for every strict transition function t the pair (t, s) is jointly stable.

Definition. We say that an arrival function is *identically distributed* if for all $A, B, C \in \mathcal{X}$ we have that $\rho(B_{\tau-1}|A_\tau) = \rho(B_{\tau-1}|C_\tau)$.

Proposition 2. *If an arrival function s is identically distributed then it is stable.*

Proposition 2 provides a partial answer to our question in terms of a sufficient condition. When an arrival function is identically distributed, in the stationary distribution, the probability that last period's choice set is B is the same for every choice set. This tells us that the probability that a choice set A inherited their preference from choice set B is the same for each A. This is sufficient to guarantee that each choice set's stationary distribution over preferences is identical.

5. Stochastic Choice

In order to capture the behavioral content of consumption dependence, in this section we look at a two period stochastic choice model where the first period is governed by RUM and choices in the second period are governed by a transition function. Formally, we observe a random joint choice rule which captures the joint choices in the first and second period.

Definition. A function $p : X^2 \times (2^X \setminus \{\varnothing\})^2 \to [0,1]$ is a random joint choice rule (rjcr) if it satisfies the following.

- $x \notin A$ or $y \notin B \implies p(x,y,A,B) = 0$
- $\sum_{x \in A, y \in B} p(x,y,A,B) = 1$

We consider two representations; one for the general case of MRUM and one for purely consumption dependent MRUM.[4]

Definition. We say that a random joint choice rule is *consistent* with MRUM if there exists a probability distribution over preferences ν and a transition function t such that the following holds for all nonempty $A, B \subseteq X$ and for all $(x,y) \in A \times B$.

$$p(x,y,A,B) = \sum_{\succ \in N(x,A)} \nu(\succ) \sum_{\succ' \in N(y,B)} t_{\succ'}(x,\succ) \tag{4}$$

Further, we say the random joint choice rule is consistent with purely consumption dependent MRUM if the transition function is purely consumption dependent.

[4]The case of purely preference dependent MRUM has been considered in Li (2022) and Chambers et al. (2022). As such, we do not consider the case of pure preference dependence here.

The concept of a Möbius inverse will be of importance for both characterizations. Consider a function $f : 2^X \setminus \{\varnothing\} \to \mathbb{R}$. The Möbius inverse of f on the partially ordered set $(2^X \setminus \{\varnothing\}, \supseteq)$ is given by the function $g : 2^X \setminus \{\varnothing\} \to \mathbb{R}$ which is recursively defined by $f(A) = \sum_{B \supseteq A} g(B)$. Each function f has a well-defined Möbius inverse which is unique to f. The Möbius inverse $g(A)$ captures how much the function f changes at A. When $g(A) \geq 0$ for each A, every set A adds some non-negative value to the function f.[5] Notably, the Block-Marschak polynomials (see Block and Marschak (1959) and Falmagne (1978)) are the Möbius inverse of choice probabilities and are used to characterize which random choice rules are consistent with RUM. The characterization asks that each Block-Marschak polynomial is non-negative. This corresponds to a non-negative amount of probability being added to the choice of each alternative at each choice set. We use multiple Möbius inverses in our upcoming characterizations.

5.1. **Pure Consumption Dependence.** In MRUM, the agent is myopic which means that they do not take into account tomorrow's choice set when making decisions. While there is still consumption dependence between the two periods as tomorrow's preference is determined by today's choice, the marginal choice of an alternative x from A in the first period should not vary with the choice set B in the second period. This leads to our first axiom.

Axiom 1. *A random joint choice rule satisfies* ***marginality*** *if for every nonempty $B, B' \subseteq X$ and for each $x \in A \subseteq X$ we have the following.*

$$\sum_{y \in B} p(x, y, A, B) = \sum_{y \in B'} p(x, y, A, B')$$

Marginality asks that the first period choice probabilities have well defined marginal choice probabilities. Marginality is not a new axiom and appears in earlier work such as Chambers et al. (2022) and in the form of marginal consistency in Li (2022). However, unlike in those two settings where marginality must hold in every period, in our setting

[5]More formally, this is related to the concept of complete monotonicity of functions on semigroups. $(2^X, \cup)$ is the semigroup we are considering here. See Berg et al. (1984) for a reference.

marginality need only hold in the first period. When a random joint choice rule satisfies marginality, we can define a marginal random joint choice rule for the first period.

Definition. For a random joint choice rule p satisfying marginality, we use $p(x, A)$ to denote the marginal choice probability of x from A in the first period.

$$p(x, A) = \sum_{y \in X} p(x, y, A, X)$$

Once we have well defined marginal choice probabilities in the first period, our model asks that these choice probabilities arise from RUM. As mentioned prior, RUM is characterized by non-negativity of the Block-Marschak polynomials which are the Möbius inverse of choice probabilities.

Definition. For each $x \in X$ and for each A such that $x \in A$, we use $q(A|x)$ to denote the Möbius inverse of $p(x, A)$.

$$\begin{aligned} q(A|x) &= p(x, A) - \sum_{A' : A \subsetneq A'} q(A'|x) \\ &= \sum_{A' : A \subseteq A'} (-1)^{|A' \setminus A|} p(x, A') \end{aligned}$$

Axiom 2. *A random joint choice rule satisfies* ***complete monotonicity of first period choice probabilities*** *if $q(A|x) \geq 0$ for each $x \in X$ and for each A with $x \in A$.*

Axioms 1-2 guarantee that the choices in the first period are consistent with RUM. All that is left is to guarantee that the second period choice probabilities arise from a purely consumption dependent transition function. As we are considering the case of pure consumption dependence, the choices of the second period are connected to the first period only through the choice in the first period and not through the preference realization in the first period. With this in mind, we consider second period conditional choice probabilities.

Definition. For each (x, A) with $p(x, A) > 0$, denote by $p(y, B|x, A)$ the choice probability of y from B in the second period conditional on x being chosen from A in the first period.

$$p(y, B|x, A) = \frac{p(x, y, A, B)}{p(x, A)}$$

Axiom 3. *We say that a random joint choice rule satisfies* ***independence of first period choice set*** *if for each (x, A) and (x, A') with $p(x, A) > 0$ and $p(x, A') > 0$, we have $p(y, B|x, A) = p(y, B|x, A')$.*

Axiom 3 asks that second period choices are dependent only on the first period choices and not on the first period choice sets. When Axiom 3 holds, we use $p(y, B|x)$ to denote $p(y, B|x, A)$. Since each purely consumption dependent transition function defines a probability distribution over preferences tomorrow conditional on today's choice, each $p(\cdot, \cdot|x)$ should be consistent with RUM.

Definition. For each x such that there exists A with $p(x, A) > 0$, for each $y \in X$, and for each B with $y \in B$, we use $q(B|x, y)$ to denote the Möbius inverse of $p(y, B|x)$.

$$\begin{aligned} q(B|x, y) &= p(y, B|x) - \sum_{B' : B \subsetneq B'} q(B'|x, y) \\ &= \sum_{B' : B \subseteq B'} (-1)^{|B' \setminus B|} p(y, B'|x) \end{aligned}$$

Axiom 4. *A random joint choice rule satisfies* ***complete monotonicity of consumption dependent second period choice probabilities****(CMCD2) if $q(B|x, y) \geq 0$ for each $x \in X$ such that there exists A with $p(x, A) > 0$, for each $y \in B$, and for each B with $y \in B$.*

Axioms 3-4 guarantee that, conditional on the choice of x in the first period, the choice probabilities in the second period are consistent with RUM and thus a purely consumption dependent transition function.

Theorem 4. *A random joint choice rule is consistent with purely consumption dependent MRUM if and only if it satisfies Axioms 1-4.*

5.2. **General Transition Functions.** In the case of pure consumption dependence, the agent's transition function only depends on their choice. We are able to condition on the agent's choice in the first period as it is an observable. However, in general, an agent's transition function will depend on both their choice and preference in the first period. This means that we will need a new set of axioms to characterize general MRUM. As the agent still chooses according to RUM in the first period, Axioms 1-2 still hold. We now introduce two new functions which will be used in the upcoming characterization.

Definition. For each $x, y \in X$ and each $A \times B$ such that $(x, y) \in A \times B$, we use $q(A|x, y, B)$ to denote the Möbius inverse of $p(x, y, A, B)$ as we vary A.

$$
\begin{aligned}
q(A|x, y, B) &= p(x, y, A, B) - \sum_{A' : A \subsetneq A'} q(A'|x, y, B) \\
&= \sum_{A' : A \subseteq A'} (-1)^{|A' \setminus A|} p(x, y, A', B)
\end{aligned}
$$

Definition. For each $x, y \in X$ and each $A \times B$ such that $(x, y) \in A \times B$, we use $r(B|x, y, A)$ to denote the Möbius inverse of $q(A|x, y, B)$ as we vary B.

$$
\begin{aligned}
r(B|x, y, A) &= q(A|x, y, B) - \sum_{B' : B \subsetneq B'} r(B'|x, y, A) \\
&= \sum_{B' : B \subseteq B'} (-1)^{|B' \setminus B|} q(A|x, y, B')
\end{aligned}
$$

Our next result offers an interpretation of these two functions. Recall that $M(x, A)$ is the set of linear orders which rank x above every element of A and below every element of $X \setminus A$.

Lemma 1. *Let (ν, t) be a MRUM representation of a consistent random joint choice rule p. Then the following holds.*

(1) $q(A|x) = \nu(M(x, A))$
(2) $q(A|x, y, B) = \sum_{\succ \in M(x,A)} \sum_{\succ' \in N(y,B)} \nu(\succ) t_{\succ'}(x, \succ)$
(3) $r(B|x, y, A) = \sum_{\succ \in M(x,A)} \sum_{\succ' \in M(y,B)} \nu(\succ) t_{\succ'}(x, \succ)$

In Lemma 1, (1) is due to Falmagne (1978) and tells us that $q(A|x)$ corresponds exactly to the probability weight put on linear orders which rank x above every element of A and below every element of $X \setminus A$. (2) tells us that $q(A|x, y, B)$ is equal to the probability of one of these prior linear orders times the probability of transitioning to a preference which chooses y from B in the second period. In other words, (2) tells us that $q(A|x, y, B)$ is the probability of choosing y from B in the second period and drawing a preference from $M(x, A)$ in the first period conditional on choosing from A in the first period. Taking this one step further, (3) tells us that $r(B|x, y, A)$ is equal to the probability of drawing a preference from $M(y, B)$ in the second period and drawing a preference from $M(x, A)$ in the first period conditional on choosing from A in the first period.

Axiom 5. *A random joint choice rule satisfies **complete monotonicity of joint choice with respect to the first period** (CMJC1) if $q(A|x, y, B) \geq 0$ for all $x, y \in X$ and for all $A, B \subseteq X$ such that $(x, y) \in A \times B$.*

In terms of MRUM, Axiom 5 tells us that, for each (x, A), the preferences in $M(x, A)$ must transition to preferences which choose y from B with non-negative probability. In terms of data, we interpret Axiom 5 in the case where Axiom 2 holds. Axiom 2 says that every choice set adds some non-negative amount to the choice probability of x. In normative terms, Axiom 5 tells us that as more probability is added to the choice of x, the joint choice probability of (x, y) must not decrease. More formally, Axiom 5 tells us that every set A contributes some non-negative probability to the joint choice of (x, y) keeping the second period choice set B fixed.

Axiom 6. *A random joint choice rule satisfies **second degree complete monotonicity of joint choice probabilities** (2DCM) if $r(B|x, y, A) \geq 0$ for all $x, y \in X$ and for all $A, B \subseteq X$ such that $(x, y) \in A \times B$.*

In terms of MRUM, Axiom 6 tells us that the transition functions $t(x, \succ)$ for $\succ \in M(x, A)$ transition to preferences in $M(y, B)$ with non-negative probability. In terms of data, we interpret Axiom 6 in the case where Axiom 2 and Axiom 5 hold. Axiom 5 tells us that while we keep the second period choice set B fixed, as the choice probability of x increases, the joint choice of (x, y) must increase. In simple terms, Axiom 6 tells us that the amount that (x, y) increases by is larger in B' than in B when $B' \subseteq B$. Alternatively, for $B' \subseteq B$, it is an implication of Axiom 6 that $q(A|x, y, B') \geq q(A|x, y, B)$. Formally, Axiom 6 tells us that each choice set B adds some non-negative probability to the amount of probability that each choice set A adds to the joint choice of (x, y). Axiom 6 can be thought of as capturing increasing differences of the joint choice probabilities while taking into account the partial order structure of set inclusion.

Theorem 5. *Consider a random joint choice rule p. The following are equivalent.*

(1) *p is consistent with MRUM.*
(2) *p satisfies Axioms 1-2 and Axioms 5-6.*
(3) *p satisfies Axiom 1 and Axiom 6.*

6. Conclusion

In this paper we discuss the impact of consumption dependence on the individual interpretation of RUM and stochastic choice. We find that while there are some forms of consumption dependence which are consistent with RUM, in general consumption dependence is not allowed by the individual interpretation of RUM. The way we model consumption dependence is Markovian, but, within the Markovian paradigm, our model is fully nonparametric. However, an analyst may not be concerned about fully nonparametric consumption dependence but may be concerned about a specific type of consumption dependence that arises in the environment they are studying. As an example, consider an analyst who is studying risk aversion and is trying to estimate the agent's level of risk aversion. In this setting, when an agent invests in a risky asset and the value of that risky asset is realized, the realization of the risky asset impacts the agent's wealth level. For general risk preferences, an agent's wealth level impacts their risk aversion. The analyst may be concerned that consumption dependence enters through this channel but is not concerned about consumption dependence entering due to habit formation. An interesting line of study would be to focus on stylized environments, like the one considered in the previous example, in order to see if specific forms of consumption dependence impact intertemporal aggregation.

In addition to studying the relation between consumption dependence and RUM, we also characterize the behavioral content of consumption dependence in a two period model.[6] Our model is characterized by two axioms. The first asks that marginal choices in the first period are independent of the choice set in the second period. Our second axiom asks that joint choice probabilities satisfy two levels of complete monotonicity and is a statement about increasing differences while taking into account the partial order structure of set inclusion. While not studied here, an interesting followup question would be to characterize static random choice rules which are consistent with MRUM. Recall that a random choice rule is consistent with MRUM if it satisfies $p(x, A) = \sum_{\succ \in N(x,A)} \nu_A(\succ)$ for some transition function t where ν_A is the stationary distribution over preferences of the Markov chain induced by t at A. If these Markov chains are not guaranteed to be ergodic, then every random choice rule can be explained if we are allowed to choose a distribution over starting choices. The interesting case is when these Markov chains are ergodic. In general, MRUM extends RUM by allowing the distribution over preferences to vary with the agent's choice set. Characterizing which static random choice rules are consistent with MRUM is of interest as MRUM may be able to explain many of the empirical irregularities of RUM such as failures of regularity.

Appendix A. Proofs

A.1. **Proof of Proposition 1.** If $t(x) = t(y)$ for all $x, y \in X$, then $t(x)$ defines the stationary distribution over preferences which is common to each choice set. Now suppose $t(x) \neq t(y)$ for some $x, y \in X$ and that t is stable with common stationary distribution ν. Consider the set $\{x, y\} = A$. With a slight abuse of notation, let $t(x, y) = \sum_{\succ \in N(x,\{x,y\})} t_\succ(x)$. Define $t(y, x)$ similarly. It then follows that t defines a Markov chain between choices on A which is given by the following matrix.

$$\begin{bmatrix} 1 - t(x, y) & t(x, y) \\ t(y, x) & 1 - t(y, x) \end{bmatrix}$$

[6]This is also done in a working paper version of Frick et al. (2019) in the context of the dynamic random expected utility model.

The first row and column of this matrix corresponds to x and the second row and column corresponds to y. The stationary distribution of this Markov chain is $\begin{bmatrix} \frac{t(y,x)}{t(y,x)+t(x,y)} \\ \frac{t(x,y)}{t(y,x)+t(x,y)} \end{bmatrix}$. Since t is strictly positive in each of its elements we can write $p(x, A) = \alpha$ and $p(y, A) = 1 - \alpha$ for some $\alpha \in (0, 1)$. It then follows that $\nu = \alpha t(x) + (1-\alpha)t(y)$. Since $t(x) \neq t(y)$, there is some $\succ$ such that, without loss of generality, $t_\succ(x) < \nu(\succ) < t_\succ(y)$. Consider $z \notin \{x, y\}$. Since the stationary distribution for each choice set is equal to ν, it must be the case that $t_\succ(x) < \nu(\succ) < t_\succ(z)$ and $t_\succ(z) < \nu(\succ) < t_\succ(y)$ which is a contradiction.

A.2. **Proof of Theorem 1.** Suppose that transition function t is stable. This means that $\nu_A = \nu_B$ for all choice sets A and B. Consider $x \in A$ with $|A| \geq 3$. Consider the following.

$$\begin{aligned} &\sum_{\succ' \in \mathcal{L}(X)} \nu_A(\succ')t_\succ(x(A, \succ'), \succ') - \sum_{\succ' \in \mathcal{L}(X)} \nu_A(\succ')t_\succ(x(A \setminus \{x\}, \succ'), \succ') \\ = &\sum_{\succ' \in N(x,A)} \nu_A(\succ')t_\succ(x, \succ') - \sum_{y \in B\setminus\{x\}} \sum_{\succ' \in N(x,A) \cap N(y, A\setminus\{x\})} \nu_A(\succ')t_\succ(y, \succ') \end{aligned}$$

The left hand side of the equation above represents $\nu_A(\succ) - \nu_{A\setminus\{x\}}(\succ)$. Since $\nu_A = \nu_{A\setminus\{x\}}$, the left hand side of the equation is equal to zero and then so too does the right hand side of the equation. Since the right hand side of the equation equals zero, a quick rearrangement gives us local stability.

Now suppose that t is locally stable. Let ν_A correspond to the stationary distribution of A used in the definition of local stability. If $|A| \geq 3$, then we have the following.

$$\begin{aligned} &\sum_{\succ' \in \mathcal{L}(X)} \nu_A(\succ')t_\succ(x(A, \succ'), \succ') - \sum_{\succ' \in \mathcal{L}(X)} \nu_A(\succ')t_\succ(x(A \setminus \{x\}, \succ'), \succ') \\ = &\sum_{\succ' \in N(x,A)} \nu_A(\succ')t_\succ(x, \succ') - \sum_{y \in B\setminus\{x\}} \sum_{\succ' \in N(x,A) \cap N(y, A\setminus\{x\})} \nu_A(\succ')t_\succ(y, \succ') \end{aligned}$$

The right hand side of the above equation is the local stability condition and thus equals zero. The left hand side of the above equation represents $\nu_A(\succ) - \langle \nu_A, m_{A\setminus\{x\}}(\succ, \cdot)\rangle$ where $\langle \cdot, \cdot \rangle$ represents the dot product of two vectors. Since the left hand side of the equation equals zero, $\nu_A(\succ) - \langle \nu_A, m_{A\setminus\{x\}}(\succ, \cdot)\rangle = 0$. This immediately gives us that $\nu_A M_{A\setminus\{x\}} = \nu_A$, and so ν_A is the stationary distribution of $A \setminus \{x\}$. By an analogous argument, for $A \neq X$, we can conclude that ν_A is the stationary distribution for $A \cup \{x\}$. Now consider some set B. We want to show that $\nu_B = \nu_A$. By prior arguments, we have shown we can either add or remove one element and maintain our stationary distribution if local stability holds. We now simply need to argue that $\nu_A = \nu_{A\cup B}$ by repeatedly making those arguments while adding a single alternative at each step. We then can show that $\nu_{A\cup B} = \nu_B$ by once again repeatedly making the same arguments while removing a single alternative at each step. Thus $\nu_A = \nu_B$ for all $A, B \in \mathcal{X}$.

A.3. **Proof of Theorem 2.** The condition for a distribution to be stationary is $\nu(M - I) = 0$ where I is the identity matrix. The condition for stability can be written similarly. Consider the matrix M with rows indexed by elements of $\mathcal{L}(X)$ and columns indexed by elements of $\mathcal{X} \times \mathcal{L}(X)$. The typical element of M is given as follows.

$$m(\succ, (A, \succ')) = t_{\succ'}(x(A, \succ), \succ) - \mathbf{1}\{\succ = \succ'\}$$

Stability can now be written as the existence of a ν such that $\nu M = 0$. Ville's Theorem of the Alternative (see Ville (1938) and Border (2013)) tells us that there exists $\nu > 0$ such that $\nu M \leq 0$ if and only if there does not exist some $b \geq 0$ such that $Mb >> 0$. Consider some $\nu > 0$ such that $\nu M \leq 0$. If such a ν exists, then we can rescale ν to be a probability distribution (i.e. $\nu \cdot 1 = 1$). M is written as a series of $(M_A - I)$ where M_A is the Markov transition matrix for choice set A in Section 2. As such $\nu M_A \cdot 1 = \nu I \cdot 1 = 1$. This means that $\nu(M_A - I) \cdot 1 = 0$ which in turn tells us that $\nu M \cdot 1 = 0$. Further, if there is some component of νM that is strictly less than zero then there must be some other component of νM strictly larger than zero. Finally, this gives us that $\nu M \leq 0$ and $\nu > 0$ if and only if $\nu M = 0$ and $\nu > 0$.

Now suppose there exists some $b \geq 0$ such that $Mb >> 0$. This b is a zero value bet from our no money pump condition. If we write out each inequality implied by

$Mb >> 0$ we get the following for each $\succ \in \mathcal{L}(X)$.

$$\sum_{(A,\succ')\in\mathcal{X}\times\mathcal{L}(X)} b(A,\succ')[t_{\succ'}(x(A,\succ),\succ) - \mathbf{1}\{\succ=\succ'\}] > 0$$

Note that $\sum_{(A,\succ')\in\mathcal{X}\times\mathcal{L}(X)} b(A,\succ')\mathbf{1}\{\succ=\succ'\}$ is exactly $b(\succ)$. This means we can rewrite the above as follows.

$$\sum_{(A,\succ')\in\mathcal{X}\times\mathcal{L}(X)} b(A,\succ')t_{\succ'}(x(A,\succ),\succ) - b(\succ) > 0$$

The negation of the above holding for all $\succ$ is exactly our no money pump condition. So our no money pump condition holds if and only if there does not exist some $b \geq 0$ with $Mb >> 0$. This is equivalent to the existence of some $\nu > 0$ satisfying $\nu \cdot 1 = 1$ and $\nu M \leq 0$ by Ville's Theorem of the Alternative. Finally we showed that this is equivalent to the existence of some $\nu > 0$ satisfying $\nu \cdot 1 = 1$ and $\nu M = 0$ which is exactly stability.

A.4. **Proof of Theorem 3.** Our proof of Theorem 3 will proceed in a similar manner to our proof of Theorem 2. Consider the matrix M with rows indexed by elements of $\mathcal{L}(X)$ and columns indexed by elements of $\mathcal{X} \times \mathcal{L}(X)$. The typical element of M is given as follows.

$$m(\succ,(A,\succ')) = \sum_{B\in\mathcal{X}} \rho(B_{\tau-1}|A_\tau)t_{\succ'}(x(B,\succ),\succ) - \mathbf{1}\{\succ=\succ'\}$$

As prior, Ville's Theorem of the Alternative tells us that there exists $\nu > 0$ such that $\nu M \leq 0$ if and only if there does not exist some $b \geq 0$ such that $Mb >> 0$. Using an analogous argument from the proof of Theorem 2, we can show that $\nu M \leq 0$ and $\nu > 0$ if and only if $\nu M = 0$ and $\nu > 0$. Now suppose there exists some $b \geq 0$ such that $Mb >> 0$. This b is a zero value bet from our no money pump condition. If we write out each inequality implied by $Mb >> 0$ we get the following for each $\succ \in \mathcal{L}(X)$.

$$\sum_{(A,\succ')\in\mathcal{X}\times\mathcal{L}(X)} b(A,\succ')[\sum_{B\in\mathcal{X}} \rho(B_{\tau-1}|A_\tau)t_{\succ'}(x(B,\succ),\succ) - \mathbf{1}\{\succ=\succ'\}] > 0$$

As before, we can rewrite the above as follows.

$$\sum_{(A,\succ')\in\mathcal{X}\times\mathcal{L}(X)}\sum_{B\in\mathcal{X}} b(A,\succ')\rho(B_{\tau-1}|A_\tau)t_{\succ'}(x(B,\succ),\succ)-b(\succ)>0$$

This is the negation of our no money pump condition. As prior this tells us that our no money pump condition holds if and only if there exists some $\nu>0$ satisfying $\nu\cdot 1=1$ and $\nu M=0$. Unlike prior, we are not done.

Our goal now is to show that the ν we just found corresponds to $\pi\nu$ in the stationary distribution of our initial Markov chain st. We now verify if $\pi\nu$ is the stationary distribution of st whenever no money pump holds.

$$\begin{aligned}
&\sum_{(\succ,B)\in\mathcal{L}(X)\times\mathcal{X}} \pi(B)\nu(\succ)t_{\succ'}(x(B,\succ),\succ)s_A(B)\\
&=\pi(A)\sum_{(\succ,B)\in\mathcal{L}(X)\times\mathcal{X}} \rho(B_{\tau-1}|A_\tau)\nu(\succ)t_{\succ'}(x(B,\succ),\succ)\\
&=\pi(A)\nu(\succ)
\end{aligned}$$

The first line above is the probability of $(\succ,A)$ in the next period given that the distribution this period is given by $\pi\nu$. The second line follows from multiplying by $\frac{\pi(A)}{\pi(A)}$ and then collecting like terms to write $\rho(B_{\tau-1}|A_\tau)$. The last line then follows from the fact that $\nu M=0$. Notably the last line holds for all $(\succ,A)$ if and only if we have $\nu M=0$ which is true if and only if our no money pump condition holds. The equality above shows that $\pi\nu$ is the stationary distribution of st if and only if no money pump holds and so we are done.

A.5. **Proof of Proposition 2.** Suppose that an arrival function s is identically distributed. Note that for any strict transition function t, we can always write the stationary distribution st as $\pi(A)\nu_A(\succ)$. Consider the following.

$$\begin{aligned}
&\sum_{(\succ,B)\in\mathcal{L}(X)\times\mathcal{X}} \pi(B)\nu_B(\succ)t_{\succ'}(x(B,\succ),\succ)s_A(B) \\
= \pi(A) &\sum_{(\succ,B)\in\mathcal{L}(X)\times\mathcal{X}} \rho(B_{\tau-1}|A_\tau)\nu_B(\succ)t_{\succ'}(x(B,\succ),\succ) \\
&= \pi(A)\nu(\succ)
\end{aligned}$$

As prior, the first line above is the probability of $(\succ, A)$ in the next period given that the distribution in this period is governed by $\pi(B)\nu_B(\succ)$. The second line follows from the same multiplication as in the proof of Theorem 3. By the fact that s is identically distributed, the summation term in the second line above is the same for every choice set A. In other words, the probability of the $\succ'$ is independent of A which gives us the third line.

A.6. **Proof of Theorem 4.** We begin with a preliminary lemma.

Lemma 2. *Let (ν, t) be a purely choice dependent MRUM representation of a consistent random joint choice rule p. When there exists (x, A) such that $p(x, A) > 0$, then $q(B|x, y) = \sum_{\succ\in M(y,B)} t_\succ(x)$*

Proof. First observe that $p(y, B|x) = \sum_{B':B\subseteq B'} q(B'|x, y)$. Now observe the following.

$$\begin{aligned}
p(y,B|x) &= \frac{p(x,y,A,B)}{p(x,A)} \\
&= \frac{\sum_{\succ'\in N(x,A)} \nu(\succ') \sum_{B':B\subseteq B'} \sum_{\succ\in M(y,B)} t_\succ(x)}{p(x,A)} \\
&= \frac{p(x,A) \sum_{B':B\subseteq B'} \sum_{\succ\in M(y,B)} t_\succ(x)}{p(x,A)} \\
&= \sum_{B':B\subseteq B'} \sum_{\succ\in M(y,B')} t_\succ(x)
\end{aligned}$$

It follows from the previous equalities that $\sum_{B':B\subseteq B'}\sum_{\succ\in M(y,B)} t_{\succ}(x) = \sum_{B':B\subseteq B'} q(B'|x,y)$. It immediately follows that $q(B|x,y) = \sum_{\succ\in M(y,B)} t_{\succ}(x)$. $\square$

We now proceed by showing necessity. Observe the following.

$$\begin{aligned}\sum_{y\in B} p(x,y,A,B) &= \sum_{\succ\in N(x,A)} \nu(\succ)\\ &= \sum_{y'\in B'} p(x,y',A,B')\end{aligned}$$

Thus Axiom 1 is necessary. From Lemma 1-(1), we know that $q(A|x) = \nu(M(x,A))$ and thus Axiom 2 is necessary as ν is a probability distribution. Observe the following.

$$\begin{aligned}p(y,B|x,A) &= \frac{\sum_{\succ\in N(x,A)} \nu(\succ)\sum_{\succ'\in N(y,B)} t_{\succ'}(x)}{p(x,A)}\\ &= \frac{\sum_{\succ\in N(x,A')} \nu(\succ)\sum_{\succ'\in N(y,B)} t_{\succ'}(x)}{p(x,A')}\\ &= p(y,B|x,A')\end{aligned}$$

The equality between the second and third lines follows from Axiom 1. This shows that Axiom 3 is necessary. By Lemma 2, we have $q(B|x,y) = \sum_{\succ\in M(y,B)} t_{\succ}(x)$ and so Axiom 4 is necessary as $t(x)$ is a probability distribution.

We now show sufficiency. By Falmagne (1978), when Axioms 1-2 hold, we know that $p(x,A)$ admits a random utility representation. Call this representation ν. Further, when Axioms 3-4 hold, Falmagne (1978) tells us that $p(y,B|x)$ has a random utility representation for each x such that $p(y,B|x)$ are defined. Call this representation $t(x)$. We now verify this representation induces our choice probabilities.

$$\begin{aligned}p(x,y,A,B) &= p(x,A)p(y,B|x)\\ &= \sum_{A':A\subseteq A'} q(A'|x) \sum_{B':B\subseteq B'} q(B'|x,y)\\ &= \sum_{\succ\in N(x,A)} \nu(\succ) \sum_{\succ'\in N(y,B)} t_{\succ'}(x)\end{aligned}$$

The equality between the second and third lines follows from our construction and Falmagne (1978). The third line is exactly our definition of consistency and so we are done.

A.7. **Proof of Lemma 1.** Lemma 1-(1) is due to Falmagne (1978) but we provide a proof here. Observe the following.

$$\sum_{A':A\subseteq A'}\sum_{\succ\in M(x,A')}\nu(\succ)=p(x,A)=\sum_{A':A\subseteq A'}q(A'|x)$$

Lemma 1-(1) follows from a standard recursive argument. We now prove Lemma 1-(2).

$$\begin{aligned}\sum_{A':A\subseteq A'}q(A'|x,y,B)&=p(x,y,A,B)\\&=\sum_{\succ\in N(x,A)}\nu(\succ)\sum_{\succ'\in N(y,B)}t_{\succ'}(x,\succ')\\&=\sum_{A':A\subseteq A'}\sum_{\succ\in M(x,A')}\nu(\succ)\sum_{\succ'\in N(y,B)}t_{\succ'}(x,\succ)\end{aligned}$$

It follows from a standard recursive argument that $q(A'|x,y,B) = \sum_{\succ\in M(x,A')}\sum_{\succ'\in N(y,B)}\nu(\succ)t_{\succ'}(x,\succ)$. This shows Lemma 1-(2) holds. We now show Lemma 1-(3).

$$\begin{aligned}\sum_{B':B\subseteq B'}r(B'|x,y,A)&=q(A|x,y,B)\\&=\sum_{\succ\in M(x,A')}\sum_{\succ'\in N(y,B)}\nu(\succ)t_{\succ'}(x,\succ)\\&=\sum_{B':B\subseteq B'}\sum_{\succ\in M(x,A')}\sum_{\succ'\in M(y,B')}\nu(\succ)t_{\succ'}(x,\succ)\end{aligned}$$

It then follows from a standard recursive argument that $r(B'|x,y,A) = \sum_{\succ\in M(x,A')}\sum_{\succ'\in M(y,B')}\nu(\succ)t_{\succ'}(x,\succ)$. Thus Lemma 1 holds.

A.8. **Proof of Theorem 5.** We will show (1) $\implies$ (3) $\implies$ (2) $\implies$ (1).

To show (1) $\implies$ (3) we need to show the necessity of Axiom 1 and Axiom 6. The necessity of Axiom 1 was shown in the proof of Theorem 4. Necessity of Axiom 6 follows from Lemma 1-(3) and the fact that ν and $t(x,\succ)$ are probability distributions.

To show (3) $\implies$ (2) recall the definition of $r(B|x,y,A)$.

$$q(A|x,y,B) = \sum_{B':B\subseteq B'} r(B'|x,y,A)$$

Axiom 6 implies that each $r(B|x,y,A)$ is non-negative. The above equation then implies that each $q(A|x,y,B)$ is non-negative and thus Axiom 5 holds. We now show that Axiom 1 and Axiom 5 imply Axiom 2.

$$\begin{aligned} q(A|x) &= \sum_{A':A\subseteq A'} (-1)^{|A'\setminus A|} p(x,A') \\ &= \sum_{A':A\subseteq A'} (-1)^{|A'\setminus A|} \sum_{y\in B} p(x,y,A',B) \\ &= \sum_{y\in B} \sum_{A':A\subseteq A'} (-1)^{|A'\setminus A|} p(x,y,A',B) \\ &= \sum_{y\in B} q(A|x,y,B) \end{aligned} \tag{5}$$

Axiom 1 implies that $q(A|x)$ and $p(x,A)$ are well defined. Axiom 5 guarantees that each $q(A|x,y,B)$ is non-negative and thus implies that each $q(A|x)$ is non-negative which is exactly Axiom 2.

We now show that (2) $\implies$ (1). From Falmagne (1978), we know that $p(x,A)$ has a random utility representation when Axioms 1-2 hold. Call this representation ν. We will construct a random utility representation for each $t(x,\succ)$ and we will choose $t(x,\succ) = t(x,\succ')$ when $\succ,\succ'\in M(x,A)$. When $q(A|x)=0$, set $t_{\succ'}(x,\succ) = \frac{1}{|\mathcal{L}(X)|}$ for each $\succ\in M(x,A)$. When $q(A|x)>0$, we use the following construction. We define a random joint choice rule $\hat{p}(y,B|x,A) = \frac{q(A|x,y,B)}{q(A|x)}$. By Equation (5), we know that $\sum_{y\in B}\hat{p}(y,B|x,A) = 1$ for each nonempty $B\subseteq X$. The Möbius inverse of $\hat{p}(y,B|x,A)$ with respect to B is given by $\frac{r(B|x,y,A)}{q(A|x)}$. By Axiom 2 and Axiom 6, $\frac{r(B|x,y,A)}{q(A|x)}$ is non-negative. By Falmagne (1978), each $\hat{p}(\cdot,\cdot|x,A)$ has a random utility representation. Call this representation $\nu_{(x,A)}$. For each $\succ\in M(x,A)$, define $t(x,\succ) = \nu_{(x,A)}$. This means that $t(x, \succ)$ is a probability distribution for each (x, A). We now verify our construction.

$$\begin{aligned}
p(x,y,A,B) &= \sum_{A':A\subseteq A'} q(A'|x,y,B) \\
&= \sum_{A':A\subseteq A'} \sum_{B':B\subseteq B'} r(B'|x,y,A') \\
&= \sum_{A':A\subseteq A', q(A'|x)>0} q(A'|x) \sum_{B':B\subseteq B'} \frac{r(B'|x,y,A')}{q(A'|x)} \\
&= \sum_{A':A\subseteq A'} \sum_{\succ\in M(x,A')} \nu(\succ) \sum_{B':B\subseteq B'} \sum_{\succ'\in M(y,B')} t_{\succ'}(x,\succ) \\
&= \sum_{\succ\in N(x,A)} \nu(\succ) \sum_{\succ'\in N(y,B)} t_{\succ'}(x,\succ)
\end{aligned}$$

The equalities in the first through third lines follow from the definitions of $q(A|x, y, B)$, $r(B|x, y, A)$, and $\hat{p}(y, B|x, A)$. The fourth line follows from our construction. The fifth line is a rearrangement. The above equalities show that our construction is consistent with the choice probabilities we began with and so we are done.

References

Afriat, S. N. (1967): "The construction of utility functions from expenditure data," *International economic review*, 8, 67–77.

Agranov, M. and P. Ortoleva (2017): "Stochastic choice and preferences for randomization," *Journal of Political Economy*, 125, 40–68.

Alós-Ferrer, C. and M. Mihm (2021): "Updating stochastic choice," *University of Zurich, Department of Economics, Working Paper.*

Arieli, I., Y. Babichenko, F. Sandomirskiy, and O. Tamuz (2021): "Feasible Joint Posterior Beliefs," *Journal of Political Economy*, 129, 2546–2594.

Barberá, S. and P. K. Pattanaik (1986): "Falmagne and the Rationalizability of Stochastic Choices in Terms of Random Orderings," *Econometrica*, 54, 707–715.

Becker, G. S. and K. M. Murphy (1988): "A theory of rational addiction," *Journal of political Economy*, 96, 675–700.

Berg, C., J. P. R. Christensen, and P. Ressel (1984): *Harmonic analysis on semigroups: theory of positive definite and related functions*, vol. 100, Springer.

Block, H. D. and J. Marschak (1959): "Random Orderings and Stochastic Theories of Response," Tech. rep., Cowles Foundation for Research in Economics, Yale University.

Border, K. (2013): "Alternative Linear Inequalities," California Institute of Technology.

Cerreia-Vioglio, S., D. Dillenberger, P. Ortoleva, and G. Riella (2019): "Deliberately Stochastic," *American Economic Review*, 109, 2425–45.

Chambers, C. P., Y. Masatlioglu, and C. Turansick (2022): "Correlated Choice," *arXiv preprint arXiv:2103.05084*.

Chen, M. K. (2008): "Rationalization and cognitive dissonance: Do choices affect or reflect preferences?" .

Chen, M. K. and J. L. Risen (2010): "How choice affects and reflects preferences: revisiting the free-choice paradigm." *Journal of personality and social psychology*, 99, 573.

Clark, S. A. (1996): "The random utility model with an infinite choice space," *Economic Theory*, 7, 179–189.

Crawford, G. S. and M. Shum (2005): "Uncertainty and learning in pharmaceutical demand," *Econometrica*, 73, 1137–1173.

Deb, R. and L. Renou (2021): "Dynamic Choices and Common Learning," *arXiv preprint arXiv:2105.03683*.

Duraj, J. (2018): "Dynamic Random Subjective Expected Utility," Unpublished.

Erdem, T. and M. P. Keane (1996): "Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets," *Marketing science*, 15, 1–20.

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

Feinberg, Y. (2000): "Characterizing common priors in the form of posteriors," *Journal of Economic Theory*, 91, 127–179.

Fiorini, S. (2004): "A Short Proof of a Theorem of Falmagne," *Journal of Mathematical Psychology*, 48, 80–82.

Frick, M., R. Iijima, and T. Strzalecki (2019): "Dynamic Random Utility," *Econometrica*, 87, 1941–2002.

Fudenberg, D. and T. Strzalecki (2015): "Dynamic Logit with Choice Aversion," *Econometrica*, 83, 651–691.

Gul, F. and W. Pesendorfer (2006): "Random expected utility," *Econometrica*, 74, 121–146.

——— (2007): "Harmful addiction," *The Review of Economic Studies*, 74, 147–172.

Heckman, J. J. (1978): "Simple statistical models for discrete panel data developed and applied to test the hypothesis of true state dependence against the hypothesis of spurious state dependence," in *Annales de l'INSEE*, JSTOR, 227–269.

——— (1981): "Heterogeneity and state dependence," in *Studies in labor markets*, University of Chicago Press, 91–140.

Kashaev, N. and V. H. Aguiar (2022): "Nonparametric Analysis of Dynamic Random Utility Models," *arXiv preprint arXiv:2204.07220*.

Lang, X. (2022): "Feasible Joint Posterior Beliefs with Many States," *Available at SSRN 4077632*.

Li, R. (2022): "An Axiomatization of Stochastic Utility," *arXiv preprint arXiv:2102.00143*.

Lu, J. and K. Saito (2020): "Repeated choice: A theory of stochastic intertemporal preferences," Tech. rep., Working paper, Social Science Working Paper, 1449. California Institute of Technology.

Machina, M. J. (1985): "Stochastic Choice Functions Generated From Deterministic Preferences Over Lotteries," *The Economic Journal*, 95, 575–594.

McFadden, D. and M. K. Richter (1990): "Stochastic Rationality and Revealed Stochastic Preference," *Preferences, Uncertainty, and Optimality, Essays in Honor of Leo Hurwicz, Westview Press: Boulder, CO*, 161–186.

Milgrom, P. and N. Stokey (1982): "Information, trade and common knowledge," *Journal of economic theory*, 26, 17–27.

Morris, S. (1994): "Trade with heterogeneous prior beliefs and asymmetric information," *Econometrica: Journal of the Econometric Society*, 1327–1347.

——— (2020): "No trade and feasible joint posterior beliefs," Tech. rep., A working paper.

Pakes, A., J. R. Porter, M. Shepard, and S. Calder-Wang (2021): "Unobserved heterogeneity, state dependence, and health plan choices," Tech. rep., National Bureau of Economic Research.

Samet, D. (1998): "Common priors and separation of convex sets," *Games and economic behavior*, 24, 172–174.

Strack, P. and D. Taubinsky (2021): "Dynamic Preference "Reversals" and Time Inconsistency," Tech. rep., National Bureau of Economic Research.

Ville, J. (1938): "Sur la théorie générale des jeux ou intervient l'habileté des joueurs," *Traité du Calcul des Probabilités et des ses Applications', Paris, Gauthiers-Villars*, 171.