Stochastic Revealed Preferences with Measurement Error: Testing for Exponential Discounting in Survey Data*

Victor H. Aguiar
vaguiar@uwo.ca

Nail Kashaev†
nkashaev@uwo.ca

This version: September 2017

Abstract Measurement error causes a loss of distributional information, preventing the researcher from applying deterministic revealed-preference tools at the individual level. This paper proposes a new statistical revealed-preference framework that is applicable to such cases. We use our framework to establish a nonparametric test for consumer models that can be characterized by their first-order conditions in environments with measurement error in consumption. The proposed testing procedure is asymptotically consistent. We provide Monte Carlo evidence showing that measurement error in the standard deterministic revealed-preference tests may lead to mistakenly judging an otherwise model-consistent consumer to be model-inconsistent. In contrast, our methodology does not reject model consistency in such cases. In the application we develop, we find support in a consumption panel survey for exponential discounting behavior in single-individual households but not in couples' households. The first finding stands in contrast with the conclusions drawn from applying the deterministic revealed-preference test of Browning (1989) to the same sample.

JEL classification numbers: C60, D10.
Keywords: rationality, quasilinear rationality, exponential discounting, time consistency, time discount factor, revealed preferences, latent variable integration via simulation, measurement error.

*We thank a coeditor and two anonymous reviewers for very constructive suggestions that substantially improved the paper. We are grateful to Roy Allen, Elizabeth Caucutt, Mark Dean, Lance Lochner, Jim MacGee, Salvador Navarro, Joris Pinkse, David Rivers, Susanne Schennach, and Todd Stinebrickner for useful comments and encouragement.

†Department of Economics, University of Western Ontario.

1. Introduction

Measurement error in consumption is a well-known feature of survey data. One important limitation of field studies that use a revealed-preference framework is that, from a statistical standpoint, they have difficulty taking into account this measurement error. This paper proposes a new statistical framework which, in the presence of measurement error, allows the formulation of a nonparametric test for any model that can be characterized by first-order conditions. Our methodology covers (but is not restricted to) the important cases of static utility maximization, quasilinear utility maximization, and the classic dynamic consumer choice problem with exponential discounting and time-separable tastes. Applying our methodology to a consumer panel data set of Spanish households, we find that dynamic rationalizability with exponential discounting cannot be rejected for single-individual households, although it is rejected for couples.

The leading solution for dealing with measurement error in the revealed-preference framework usually consists of perturbing (minimally) any observed individual consumption streams in order to satisfy the conditions of a revealed-preference test (Adams et al. (2014)). However, this approach does not allow for standard statistical hypothesis testing. In particular, one cannot control the probability of making an error in rejecting a particular model when such a rejection could be an artifact of noisy measurements. A key concern is that a deterministic revealed-preference test may overreject the null hypothesis in the presence of measurement error. We provide Monte Carlo evidence that this concern may be very relevant in practice: the standard deterministic revealed-preference test can reject rationalizable behavior with high probability (close to 1) due to perturbations caused by only moderate measurement error.

Our methodology for taking measurement error into account is fully nonparametric. We require only a weak centering assumption on the unobserved consumption measurement error. We need that in a given decision problem, on average, consumers be accurate in recalling and reporting their total expenditures. Our approach takes advantage of the work of Schennach (2014) on Entropic Latent Variable Integration via Simulation (ELVIS). We use our main result to formulate a statistical test for the null hypothesis of the rationalizability of a random array of prices and consumption by any model that can be characterized by first-order conditions.

We wish to emphasize that our statistical methodology can be applied to several other revealed-preference models including (but not limited to) firm cost minimization (Varian (1984)), dynamic rationalizability with quasi-hyperbolic discounting (Blow et al. (2017)), homothetic rationalizability (Varian (1985)), and static utility maximization with nonlinear budget constraints (Forges and Minelli (2009)). However, for the sake of concreteness and due to their importance for applied work, we focus on the static utility maximization model (Afriat (1967)), the quasilinear utility maximization model (Brown and Calsamiglia (2007)), and the intertemporal consumption model with exponential discounting behavior (Browning (1989)).

We also provide a general methodology to make out-of-sample predictions or counterfactual analysis with minimal assumptions. In fact, we do not need to make parametric assumptions about preferences or heterogeneity, nor to impose strong distributional assumptions on measurement error.[1]

Our *empirical contribution* is to apply our methodology to a consumer panel data set gathered from single-individual and couples' households in Spain to test for dynamic rationalizability with exponential discounting.[2] Under the exponential discounting model, the consumer's time preferences are captured by a time-invariant discount factor and a time-invariant instantaneous utility. The main feature of this model is the time-consistency of the exponential discounting consumer. In other words, if the consumer prefers consumption bundle c at time t to x at time $t+k$, then she will always prefer c at time τ to x at time $\tau + k$.

The exponential discounting model remains the workhorse of a large body of applied work in economics. Nevertheless, this model is under increasing scrutiny in the light of experimental evidence that tends to find that the behavior of experimental subjects is time-inconsistent. Unfortunately, there is less evidence from the field concerning time-inconsistent behavior, and in particular, from survey data. Several authors, such as Browning (1989), DellaVigna and Malmendier (2006), and Blow et al. (2017), have provided suggestive evidence from the field against the validity of the exponential discounting model. Our methodology addresses, in a nonparametric fashion, the presence of measurement error in survey-gathered data, in order to examine the robustness of these findings from the field.

We find support for exponential discounting behavior for single-individual households. This is in contrast to the results of applying the deterministic methodology of Browning (1989) to the same sample. At the same time, in line with the findings of Blow et al. (2017) (who also use Browning (1989)'s deterministic methodology), we reject the null hypothesis of exponential discounting for the case of couples. When compared with the single-household evidence, these results suggest that time inconsistencies in consumption behavior in the couples' case arise due to preference aggregation. Hence, we provide evidence that not all couples' households display exponential discounting behavior, even when we account for measurement error. We also examine a generalized version of the exponential discounting model in the form of the collective household consumer unit proposed by Adams et al. (2014). In this more general model, each individual within the couple is an exponential discounter, and the individuals jointly decide how much to consume privately and publicly by maximizing a weighted sum of individual discounted utilities (under an assumption that requires full efficiency in the provision of publicly consumed goods). We fail to reject this particular version of the collective model with exponential discounting behavior for the couples case.

[1] Marschak (1974) suggested that a counterfactual exercise should start by establishing the minimal set of assumptions and data needed to make predictions of interest. We believe that, by combining the first-order conditions revealed-preference approach with ELVIS, our methodology satisfies this principle.

[2] This data set has been used in Beatty and Crawford (2011), Blow et al. (2013), and Adams et al. (2014).

Outline

The paper proceeds as follows. Section 2 presents a brief literature review about previous studies to extend the revealed-preference methodology to the case of measurement error. Section 3 presents the first-order conditions approach to the deterministic revealed-preference methodology that we use as a benchmark. Section 4 contains the main contribution: our new statistical nonparametric test for any model characterized by its first-order conditions under shape constraints in the presence of measurement error. Here, we introduce a centering condition on the measurement error on consumption. We formulate a new statistical notion of rationalizability on the basis of the ELVIS methodology (Schennach (2014)). Section 5 establishes a testing and inference framework for statistical rationalizable behavior. Section 6 implements our empirical test for the case of dynamic rationalizability with exponential discounting in a well-known consumer panel-survey data set used in Adams et al. (2014). Here we report the results for the testing procedure for the single-individual households, and for the couples' households. Section 7 presents a new statistical framework for recoverability and out-of-sample predictions on the basis of our testing methodology. Finally, we conclude in Section 8. All proofs can be found in the appendix. The appendix also contains two Monte Carlo experiments that assess the performance of our testing procedure in finite samples. In addition, the appendix includes a variant of our methodology for dealing with measurement error in both prices and consumption, and an application of our test to the collective exponential discounting model.

2. Literature Review

In this work, we exploit the fact that many models of consumer choice (and of decision-making in general), both static and dynamic in nature, can be fully characterized by their first-order conditions. In his seminal work, Afriat (1967) shows that in order to test the consistency of a finite data set of prices and consumption with a model of interest, it is sufficient to impose shape constraints on the unobserved utility function and thereby bypass the need to parametrize such first-order conditions. We generalize this insight by allowing measurement error in consumption. Among authors using the deterministic revealed-preference approach, the immediate antecedents to our work using the first-order approach are (i) Browning (1989) in the context of dynamic rationalizability with exponential discounting, (ii) Blow et al. (2017) for dynamic rationalizability with quasi-hyperbolic discounting, and (iii) Brown and Calsamiglia (2007) for static quasilinear utility maximization. (All these cases are applications of Rockafellar (1970)). Important advances have been made on testing and doing counterfactual analysis under random rationalizability or random utility.[3] However, the majority of these results assume that observed quantities are measured accurately. In this regard,

[3]Relevant examples are Blundell et al. (2014), Dette et al. (2016), Kitamura and Stoye (2016), and Lewbel and Pendakur (2017).

these works are focused mostly on the heterogeneity of preferences. In contrast, our framework is concerned with both heterogeneity and measurement error.

Varian (1985) is possibly the first work to introduce the subject of measurement error into the revealed-preference approach. Varian's methodology is the closest to that of our own work; he considers precisely measured (albeit random) prices to study measurement error in consumption. Varian's work is compatible with standard statistical hypothesis testing under the strong assumptions of normality (with known variance) and additivity of consumption measurement error. In contrast, our methodology is fully nonparametric. We are able to improve upon Varian's methodology and relax its core assumptions by using the nonparametric approach of Schennach (2014) to deal with measurement error.

Other papers have dealt with measurement error under different parametric assumptions about measurement error or about the heterogeneity of preferences. Gross (1995) assumes that random consumption is generated by consumers with similar preferences. In contrast, our approach allows for unrestricted heterogeneity in preferences. Tsur (1989) imposes a log-normal multiplicative measurement-error structure in expenditures. Hjertstrand (2013) proposes a generalization of Tsur (1989) and of Varian (1985), one which provides a statistical framework for testing revealed-preference consistency but requires knowing the distribution of measurement error. Our approach only requires a centering condition on the measurement error distribution. In a recent paper, Deb et al. (2017) consider a nonparametric model of "price preference." They propose a revealed-preference test of their model (based on Kitamura and Stoye (2016)) that is robust to small measurement error in prices (their main focus is on sampling error when testing using repeated cross sectional data in a random utility framework). In contrast, our methodology allows for substantially large perturbations of observed quantities. Boccardi (2016) considers a case of demand with error (possibly measurement error) and focuses on establishing a way to account for the trade-off between the fit of the model and its predictive ability. In her framework, rationality loses its empirical content and statistical hypothesis testing is not possible (which is a generalization of Beatty and Crawford (2011)).

In practice, the revealed-preference theorists (e.g., Adams et al. (2014), and Cherchye et al. (2017)) have dealt with measurement error by perturbing (minimally) the observed individual consumption in order to satisfy the conditions of a revealed-preference test. For instance, Adams et al. (2014) find the additive perturbation with a minimal norm that renders the individual consumption streams compatible with the revealed-preference restrictions. Then a subjective threshold is imposed on the maximum admissible norm of the measurement error vector. If the computed norm is above the threshold, then the model is rejected. However, their methodology has one important drawback: every data set can be made to satisfy their test or, equivalently, the test has no power (given the subjectivity of the threshold). In addition, it is unclear in their model what the probability of rejecting the null hypothesis is when the null hypothesis is in fact true. This is undesirable for a statistical test. In contrast, our approach allows one to perform traditional statistical hypothesis testing without compromising the generality of the revealed-preference conditions under a weak centering assumption on measurement error. This is important because in our methodology the critical value used to conduct the test arises from the asymptotic behavior of our test statistic.[4]

Among researchers using the revealed-preference approach, Blundell et al. (2003) are the first to provide consumer demand bounds, under the assumption of static utility maximization in a semiparametric environment (with additive noise) in which income changes continuously. Our work differs from theirs mainly in that we do not require that income be observable, nor we impose semiparametric assumptions on wealth effects to provide bounds for demand, given new prices. In addition, since they do not focus on measurement error in consumption, it is not clear if nonclassical measurement error can be compatible with their approach.[5]

We apply our methodology to a consumer panel data set gathered from single-individual and couples' households in Spain to test for dynamic rationalizability with exponential discounting. This important model is under increasing scrutiny because experimental evidence tends to find that the behavior of experimental subjects is time-inconsistent. Andreoni and Sprenger (2012) reject the null hypothesis of exponential discounting in favor of hyperbolic discounting (present bias). In an experimental pilot study, Montiel Olea and Strzalecki (2014) set-identify the distribution of hyperbolic discounting and find evidence against exponential discounting. More recently, Echenique et al. (2014) develop a revealed-preference methodology (for one good) with several consumption streams (nonparametric and deterministic). They use the data set from Andreoni and Sprenger (2012) to conclude that the majority of subjects are not consistent with exponential discounting. Nonetheless, it is important to explore to what extent their finding has external validity.

To address this issue, some researchers have turned to survey data in the form of household consumption panels. Most of this survey work has found evidence against exponential discounting. However, the existing literature has not yet addressed the issue of measurement error in the consumption reported by households in a way that allows us to perform traditional hypothesis testing. (Some additional problems with the existing evidence are (i) the strong parametric assumption on preferences, and (ii) homogeneity restrictions on the discount factor and preferences.)

One solution to some of the problems in the literature can be found in the work on deterministic revealed preference by Browning (1989). In particular, Browning's work avoids making parametric assumptions about the functional form of instantaneous utility. However, this work does not take into consideration the fact that consumption quantities can be mismeasured. Blow et al. (2017) apply the deterministic methodology of Browning (1989) to survey data (similar to the data set we use) to find that less than 1 percent of the couples' households pass the revealed-preference test for exponential discounting. In addition, when we applied their deterministic methodology to our single-individual households data set, we also obtained a very low percentage success rate. However, this low success rate of the deterministic test for exponential discounting may be due to measurement error. In our empirical application, we found support for exponential discounting behavior in single-individual households, while at the same time, support for the negative finding in

[4] A recent paper by Gillen et al. (2017) proposes parametric econometric techniques to deal with measurement error in experimental setups. They find that several experimental findings may not be robust to the presence of measurement error.

[5] Weyl (2009) discusses a parametric methodological approach to identify (not to test) the first-order conditions. Our work differs from that interesting approach in that we do not require point identification. Instead, we focus on set-valued predictions in a fully nonparametric framework.

Blow et al. (2017) in the case of couples' households. Moreover, in our Monte Carlo experiment, the deterministic test in Browning (1989) rejects the correct null hypothesis of exponential discounting behavior in 61.5 percent and 62.3 percent of the cases on average across 1000 trials with samples of size 200 or 1500, respectively, while our methodology correctly accepts the null hypothesis that all single households are consistent with random exponential discounting. (See section 10.2.)

The work of DellaVigna and Malmendier (2006) focuses on preferences over contracts, and time-discounting behavior in the field. This study uses a parametric utility specification and the assumption that time discount factors are common across individuals. They allow for measurement error with parametric restrictions. Under these assumptions they find evidence against the exponential discounting model. More importantly, their work differs from our empirical application in the following respect. Their work is a quasi-experimental study that uses external sources of variation to differentiate between the exponential and hyperbolic cases. In contrast, our work deals with survey data sets in which there is no other information beyond interest rates, prices, and consumption streams. In this environment, we establish that we can differentiate exponential discounting behavior from sophisticated hyperbolic discounting behavior. However, we do not reject the null hypothesis of exponential discounting for the case of single-individual households. We note that consumers may be hyperbolic discounters in one-shot decisions (e.g., signing a gym contract), but may be time-consistent in making recurring decisions (e.g., buying food, recreation, or transportation), as in our application.

Our empirical application also contributes to the literature on estimating the discount factor distribution in survey data sets and in a classical consumer theory environment. This has been the topic of a large body of work which, however, has reached little or no consensus.[6] This lack of consensus can be attributed in some degree to a failure to identify the parameters of interest. Here, we show that the discount factor distribution cannot be identified solely from prices, interest rates, and consumption observations in a data set that suffers from measurement error. (For details see section 7.) In this situation, only the support of the distribution can be set-identified. However, the methodology presented in this paper allows us to test for exponential discounting behavior even in this setting (i.e., without identifying the discount factor distribution).[7]

If one ignores the issues of measurement error, the Euler equation allows one to estimate the discount factor and the marginal utility either parametrically or semi-parametrically.[8] Since our objective is not to estimate but to test the exponential discounting model, we follow a different path. We have focused on ways to eliminate the latent infinite-dimensional parameters (e.g., the utility functions) by exploiting their shape restrictions and first-order conditions.

[6]We refer the reader to the survey by Frederick et al. (2002), for its extensive references, since we focus here only on the immediate antecedents to our work.

[7]In order to learn more information about the discount factor distribution, one needs additional data. One notable example is Mastrobuoni and Rivers (2016), which uses a quasi-experiment to pin down criminals' time preferences.

[8]Examples of estimators of the Euler equation and similar models include Hall (1978), Hansen and Singleton (1982), Dunn and Singleton (1986), Gallant and Tauchen (1989), Chapman (1997), Campbell and Cochrane (1999), Ai and Chen (2003), Chen and Ludvigson (2009), Darolles et al. (2011), Chen et al. (2014), and Escanciano et al. (2016).

3. The Revealed-Preference Methodology and the First-Order Conditions Approach

The main objective of this section is to provide a brief summary in a united fashion of three very important deterministic consumer models and their revealed-preference characterization. In particular, we study the models of static utility maximization (R) and quasilinear utility maximization (Q). We also investigate the case of dynamic rationality, also called the exponential discounting (ED) model. These models are at the center of many applied and theoretical works. We show that they can be completely characterized by their first-order conditions in a revealed-preference fashion. All quantities used here are assumed to be measured precisely.

Let the consumption space be $\mathbb{R}^L_+ \setminus \{0\}$, where $L \in \mathbb{N}$ is the number of commodities.[9] Consider a consumer who is endowed with a utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$ that is assumed to be concave, locally nonsatiated, and continuous. The consumer faces a sequence of decision problems indexed by $t \in \mathcal{T}$, where $\mathcal{T} = \{0, \cdots, T\}$, with a known and finite $T \in \mathbb{N}$. At each decision problem $t \in \mathcal{T}$, the consumer faces the price vector $p_t \in \mathbb{R}^L_{++}$.

Definition 1. (Static m-rationalizability) For any $\mathrm{m} \in \{\mathrm{R}, \mathrm{Q}\}$, a deterministic array $(p_t, c_t)_{t\in\mathcal{T}}$ is m-rationalizable (in a static sense) if for some constants $z^{\mathrm{m}}_t \in \mathbb{R}$ and $y_t > 0$, $t \in \mathcal{T}$, the consumption bundle c_t solves:

$$\begin{aligned} &\max_{c\in\mathbb{R}^L_+} u(c) + z^{\mathrm{m}}_t, \\ &\text{s.t. } p_t'c + z^{\mathrm{m}}_t = y_t, \end{aligned}$$

for all $t \in \mathcal{T}$.

The rational case requires that the constant $z^{\mathrm{R}}_t = 0$ for all $t \in \mathcal{T}$. The quasilinear utility case requires that $z^{\mathrm{Q}}_t > 0$ for all $t \in \mathcal{T}$. The interpretation of z^{Q}_t is that it is a numeraire or composite outside commodity.

Next we focus our attention on the dynamic consumer model of exponential discounting (ED). We assume that an individual consumer has preferences over a stream of dated consumption bundles $(c_t)_{t\in\mathcal{T}}$, where $\mathcal{T} = \{0, \cdots, T\}$, $T \in \mathbb{N}$, and $c_t \in \mathbb{R}^L_+ \setminus \{0\}$. (The number of goods, L, is kept the same across the time interval.) At time τ, the consumer chooses how much c_τ she will consume by maximizing

$$V_\tau(c) = u(c_\tau) + \sum_{j=1}^{T-\tau} d^j u(c_{\tau+j}),$$

[9] We use $\mathbb{N}$ to denote the set of natural numbers. The expression $\mathbb{R}^L_+$ denotes the set of componentwise nonnegative elements of the L-dimensional Euclidean space $\mathbb{R}^L$, and $\mathbb{R}^L_+ \setminus \{0\}$ denotes the set of vectors $v \in \mathbb{R}^L_+$ that are distinct from zero ($v \neq 0$). Similarly, $\mathbb{R}^L_{++}$ denotes the set of componentwise positive elements of $\mathbb{R}^L_+$. The inner product of two vectors $v_1, v_2 \in \mathbb{R}^L$ is denoted by $v_1'v_2$.

subject to the linear budget or flow constraints shown here:

$$p_t'c_t - y_t + s_t - a_t = 0, \quad t = \tau, \ldots, T,$$

where $d \in (0, 1]$ is the discount factor; $p_t \in \mathbb{R}^L_{++}$ is the price vector as before; $y_t \in \mathbb{R}_{++}$ is income received by the individual at time t; s_t is the amount of savings held by the consumer at the end of time t; and a_t is the volume of assets held at the start of time t. The consumer invests all her savings. Moreover, the assets evolve according to the following law of motion:

$$a_t = (1 + r_t)s_{t-1},$$

where $r_{t+1} > -1$ is the interest rate that is accessible for the consumer. The holdings of assets in the last period ($t = T$) are set to be zero.

The intertemporal value function, $V_t : \mathbb{R}^{L\times(T-t+1)}_{+} \to \mathbb{R}_{++}$, represents the consumer preferences at a given time t. The components of this representation are the parameters of the model. First, $d \in (0, 1]$ is a scalar number that measures the degree of discount that the consumer gives to the future. Second, $u : \mathbb{R}^L_+ \to \mathbb{R}_{++}$ is an instantaneous utility function that is assumed to be concave, locally nonsatiated, and continuous. The exponential discounting consumer is time-consistent, that is, she will solve the dynamic problem above the same way at any point of the time window. In particular, the time-consistent consumer will be able to keep her commitment to the solution to the problem at the first time period ($\tau = 0$).

Definition 2. (ED-rationalizability) A deterministic array $(p_t, r_t, c_t)_{t\in\mathcal{T}}$ is ED-rationalizable if there exist $(y_t)_{t\in\mathcal{T}} \in \mathbb{R}^{|\mathcal{T}|}_{++}$ and $a_0 \geq 0$ such that the consumption stream $(c_t)_{t\in\mathcal{T}}$ solves:

$$\max_{z\in\mathbb{R}^{L\times|\mathcal{T}|}_{+}} u(z_0) + \sum_{t=1}^{T} d^t u(z_t),$$

subject to

$$p_0'z_0 + \sum_{t=1}^{T} \frac{p_t'z_t}{\prod_{i=1}^{t}[1 + r_i]} = \sum_{t=1}^{T} \frac{y_t}{\prod_{i=1}^{t}[1 + r_i]} + a_0.$$

3.1. The First-Order Conditions Approach

Now we establish that any consumer model $\mathrm{m} \in \{\mathrm{R}, \mathrm{Q}, \mathrm{ED}\}$ can be completely characterized in terms of its first-order conditions with respect to (i) a concave, locally nonsatiated, and continuous utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$, (ii) the effective (or transformed) prices $\rho^{\mathrm{m}}_t \in \mathbb{R}^L_{++}$, and (iii) restrictions on some constants $\lambda^{\mathrm{m}}_t \in \mathbb{R}_{++}$, interpreted as the marginal utility of income, for each $t \in \mathcal{T}$. We call this *the first-order conditions approach.* Observe that the utility function is model-independent, but the effective prices and marginal utility of income are not. We define the effective prices in Table 1.

Table 1 – Definition of ρ_t^{m}

m	R	Q	ED
ρ_t^{m}	p_t	p_t	$p_t/\prod_{j=1}^{t}(1+r_j)$

Let $\nabla u(c_t)$ denote a supergradient of u at the point c_t. (Under differentiability, $\nabla u(c_t)$ is a gradient.)

Lemma 1. *For any model* $\mathrm{m} \in \{\mathrm{R}, \mathrm{Q}, \mathrm{ED}\}$*, a deterministic array* $(\rho_t^{\mathrm{m}}, c_t)_{t\in\mathcal{T}}$ *is* m*-rationalizable if and only if there exists a pair* $(u, (\lambda_t^{\mathrm{m}})_{t\in\mathcal{T}})$ *such that*

(i) $u : \mathbb{R}_+^L \to \mathbb{R}$ *is a concave, locally nonsatiated, and continuous utility function;*

(ii) $\nabla u(c_t) \leq \lambda_t^{\mathrm{m}}\rho_t^{\mathrm{m}}$ *for every* $t \in \mathcal{T}$*. If* $c_{t,j} \neq 0$*, then* $\nabla u(c_t)_j = \lambda_t^{\mathrm{m}}\rho_{t,j}^{\mathrm{m}}$*, where* $c_{t,j}$*,* $\nabla u(c_t)_j$*, and* $\rho_{t,j}$ *are the* j*-th components of* c_t*,* $\nabla u(c_t)$*, and* ρ_t*, respectively.*

(iii) $\lambda_t^{\mathrm{R}} = \lambda_t > 0$*,* $\lambda_t^{\mathrm{Q}} = 1$*, and* $\lambda_t^{\mathrm{ED}} = d^{-t}$*, where* $d \in (0, 1]$*, for all* $t \in \mathcal{T}$*.*

This lemma summarizes the results in Brown and Calsamiglia (2007) for the quasilinear case and those in Browning (1989) for the exponential discounting case, and it is trivial for the static rationalizability case. Even if we focus on these three models for expositional and motivational purposes, our methodology is applicable to any model that can be characterized using the first-order conditions approach.

Remark 1. Lemma 1 allows for nondifferentiable utility functions. So, the supergradient of $u(c_t)$ may be set-valued. In this case one should read the condition $\nabla u(c_t) \leq \lambda_t^{\mathrm{m}}\rho_t^{\mathrm{m}}$ as "there exists $\xi \in \nabla u(c_t)$ such that $\xi \leq \lambda_t^{\mathrm{m}}\rho_t^{\mathrm{m}}$."

3.2. The Elimination of a Latent Infinite-Dimensional Parameter

Since our objective is not to estimate but to test m-rationalizability, we will eliminate the utility function u from its characterization. We follow the theorists of revealed-preference to eliminate the latent infinite-dimensional parameters by exploiting their shape restrictions.

In particular, we follow Afriat (1967), Varian (1985), Browning (1989), and Rockafellar (1970) to formulate a result that eliminates the utility function u (an infinite dimensional parameter) from the first-order conditions. The cost of doing this is that we have to replace the first-order conditions by a set of inequalities. However, these inequalities do not depend any more on the infinite-dimensional parameter u. Moreover, they require only the concavity of u. As a result, the inequalities are exact and do not involve any form of approximation; this is an advantage compared to other nonparametric methods (e.g., sieves, kernel estimators) or the parametric approach used in many applied papers.

To formulate our result, we first recall the definition of the concavity of u.

Definition 3. (Concavity) A utility function u is said to be concave if and only if $u(c_s) - u(c_t) \leq \nabla u(c_t)'(c_s - c_t)$, for all $s, t \in \mathcal{T}$.

Remark 2. In Definition 3 we implicitly assume the existence of the supergradient of u. Since the supergradient may be set-valued, one should read the condition $u(c_s) - u(c_t) \leq \nabla u(c_t)'(c_s - c_t)$ as "$u(c_s) - u(c_t) \leq \xi'(c_s - c_t)$ for all $\xi \in \nabla u(c_t)$."

The nonparametric characterization of the m-rationalizability of (i) observed consumption and (ii) prices without measurement error is captured by the following result.

Theorem 1. *For any $m \in \{\mathrm{R}, \mathrm{Q}, \mathrm{ED}\}$, the following are equivalent:*

(i) The deterministic array $(\rho_t^{\mathrm{m}}, c_t)_{t\in\mathcal{T}}$ is m*-rationalizable.*

(ii) There exist a vector $(\lambda_t^{\mathrm{m}})_{t\in\mathcal{T}}$ and a positive vector $(v_t)_{t\in\mathcal{T}}$ such that:

$$v_t - v_s \geq \lambda_t^{\mathrm{m}} \rho_t^{\mathrm{m}\prime}(c_t - c_s),$$

with $\lambda_t^{\mathrm{R}} = \lambda_t > 0$, $\lambda_t^{\mathrm{Q}} = 1$, and $\lambda_t^{\mathrm{ED}} = d^{-t}$, where $d \in (0, 1]$, for all $t, s \in \mathcal{T}$.

Theorem 1 summarizes known results from the revealed-preference literature.[10] Observe that Theorem 1 has transformed the first-order conditions that depend on the infinite-dimensional u to a set of inequality conditions that depend only on a deterministic finite-dimensional array $(v_t, \lambda_t^{\mathrm{m}})_{t\in\mathcal{T}}$. Nonetheless, this set of conditions is satisfied if and only if we can find a utility function that satisfies the conditions in Lemma 1. Checking the set of inequalities is (usually) a parametric linear programming problem that tells us whether a consumption stream is m-rationalizable.

This methodology is traditionally applied at the individual level in survey data, assuming that the data contains no measurement error. The results are often disappointing, with high rates of rejections for some of the models of interest. We argue that this may be an overly pessimistic conclusion, and possibly an artifact of measurement error. For that reason, in the next section we extend the revealed-preference framework to a noisy or stochastic environment. (Testing our version of statistical m-rationalizability is a convex programming problem.) We illustrate this point in the application section (section 6) by testing for ED-rationalizability in the presence of measurement error.

4. The Revealed-Preference Approach with Measurement Error

In this section, we introduce a new statistical notion of m-rationalizability (henceforth, s/m-rationalizability) with mismeasured consumption, and provide a result similar to Theorem 1 in the presence of measurement error.

[10] The proof is a consequence from the results in Afriat (1967), Varian (1985), and Browning (1989) taken together.

4.1. Statistical Rationalizability

We are interested in testing a statistical model of consumption such that each individual is an independent, identically distributed (i.i.d.) draw from some stochastic consumption rule. Using Lemma 1 as motivation, we directly define s/m-rationalizability as follows. From here on, we use boldface font to denote random objects and regular font for deterministic ones. Let $\boldsymbol{\rho}_t \in P_t \subseteq \mathbb{R}^L_{++}$ and $\mathbf{c}^*_t \in C^*_t \subseteq \mathbb{R}^L_+ \setminus \{0\}$ denote random vectors of effective prices and true consumption at time t, respectively.[11]

Definition 4. (s/m-rationalizability) A random array $(\boldsymbol{\rho}_t, \mathbf{c}^*_t)_{t\in\mathcal{T}}$ is s/m-rationalizable if there exists a pair $(\mathbf{u}, (\boldsymbol{\lambda}_t)_{t\in\mathcal{T}})$ such that

(i) $\mathbf{u}$ is a random, concave, locally nonsatiated, and continuous utility function;

(ii) $(\boldsymbol{\lambda_t})_{t\in\mathcal{T}}$ is a positive random vector, interpreted as the marginal utility of income, supported on or inside a known set $\Lambda \subseteq \mathbb{R}^{|\mathcal{T}|}_{++}$;

(iii) $\nabla \boldsymbol{u}(\mathbf{c}^*_t) \leq \boldsymbol{\lambda_t}\boldsymbol{\rho}_t$ a.s. for all $t \in \mathcal{T}$;

(iv) For every $j = 1, \ldots, L$ and $t \in \mathcal{T}$, it must be that $\mathbb{P}\left(\mathbf{c}^*_{t,j} \neq 0, \nabla \boldsymbol{u}(\mathbf{c}^*_t)_j < \boldsymbol{\lambda_t}\boldsymbol{\rho}_{t,j}\right) = 0$, where $c^*_{t,j}$, $\rho_{t,j}$, and $\nabla u(c^*_t)_j$ denote the j-th components of c^*_t, ρ_t, and $\nabla u(c^*_t)$, respectively.

This definition means that for a given realization of (i) the utility function and (ii) the marginal utility of income, the realized effective prices and the realized true consumption should fulfill the inequality $\nabla u(c^*_t) \leq \lambda_t \rho_t$.[12] This is a special case of the dynamic random utility model in which the preferences (captured by $\mathbf{u}$) and the distribution of the marginal utility of income (captured by $(\boldsymbol{\lambda})_{t\in\mathcal{T}}$) are drawn at some initial time for each consumer, and then are kept fixed over time.

Several consumer models can be characterized by their first-order conditions and by restrictions on the marginal utility of income, as we observed in section 3. For instance, we define the statistical version of R-rationalizability or s/R-rationalizability by requiring that the support of the marginal utility of income be strictly positive (i.e., $\Lambda = \mathbb{R}^{|\mathcal{T}|}_{++}$). Similarly, we define s/Q-rationalizability by requiring that the support Λ of $(\boldsymbol{\lambda}_t)_{t\in\mathcal{T}}$ be given by the restriction $\boldsymbol{\lambda}_t = 1$ a.s. for all $t \in \mathcal{T}$. Finally, we define s/ED-rationalizability by imposing that the support Λ be given by the restriction $\boldsymbol{\lambda}_t = \mathbf{d}^{-t}$, where $\mathbf{d}$ is supported on $(0, 1]$. The effective prices in each case have to be defined according to Table 1. Henceforth, we fix some model such that the effective prices $(\boldsymbol{\rho}_t)_{t\in\mathcal{T}}$ and the support Λ are known.

[11] For short, we use a.s. instead of "almost surely." We denote (i) the probability of an event A by the expression $\mathbb{P}(A)$; (ii) the indicator function by $\mathbb{1}(A) = 1$ when the statement A is true, otherwise it is zero; (iii) the mathematical expectation of any random vector $\mathbf{z}$ by the expression $\mathbb{E}[\mathbf{z}]$; (iv) the condition expectation of any random vector $\mathbf{z_1}$ conditional on another random vector $\mathbf{z_2}$ by the expression $\mathbb{E}[\mathbf{z_1}|\mathbf{z_2}]$; (v) the cardinality of a set $\mathcal{A}$ is given by the expression $|\mathcal{A}|$; (vi) the norm of a vector v is given by $\|v\|$; and (vii) the independence of two random variables $(\mathbf{z_1}, \mathbf{z_1})$ by the expression $\mathbf{z_1} \perp\!\!\!\perp \mathbf{z_2}$.

[12] The definition of the s/m-rationalizability allows for nondifferentiable utility functions. So, the supergradient of $\boldsymbol{u}(\mathbf{c}^*_t)$ may be set-valued. In this case one should read the condition $\nabla u(c^*_t) \leq \lambda_t \rho_t$ as "there exists $\xi \in \nabla u(c^*_t)$ such that $\xi \leq \lambda_t \rho_t$."

Given the definition of s/m-rationalizability, we can now formulate the stochastic version of Theorem 1.

Lemma 2. *For a given random array $(\boldsymbol{\rho}_t, \mathbf{c}_t^*)_{t\in\mathcal{T}}$, the following are equivalent:*

(i) The random array $(\boldsymbol{\rho}_t, \mathbf{c}_t^*)_{t\in\mathcal{T}}$ is s/m*-rationalizable.*

(ii) There exist positive random vectors $(\mathbf{v}_t)_{t\in\mathcal{T}}$ and $(\boldsymbol{\lambda_t})_{t\in\mathcal{T}}$ supported on or inside $\Lambda \subseteq \mathbb{R}_{++}^{|\mathcal{T}|}$ such that

$$\mathbf{v}_t - \mathbf{v}_s \geq \boldsymbol{\lambda}_t \boldsymbol{\rho}_t'(\mathbf{c}_t^* - \mathbf{c}_s^*) \quad \text{a.s.}, \quad \forall s, t \in \mathcal{T}.$$

Lemma 2 allows us to statistically test the s/m-rationalizability of $(\boldsymbol{\rho}_t, \mathbf{c}_t^*)_{t\in\mathcal{T}}$. However, any test based on this notion of rationalizability cannot differentiate between quasi-consistent cases and exact s/m-rationalizabilty (an issue first identified by Galichon and Henry (2013)). The reason is that there can be random choice rules that are not m-rationalizable, and that are arbitrarily close to being s/m-rationalizable. The issue arises because the set of s/m-rationalizable behaviors may not be closed. That is why we need to extend the notion of the consistency of a data set that is characterized by s/m-rationalizability.

Example 1. (Hyperbolic Discounting) Consider the case of a consumer who maximizes

$$V_\tau(c) = u(c_\tau) + \beta \sum_{j=1}^{T-\tau} d^j u(c_{\tau+j}),$$

where $\beta \in (0, 1]$ is the present-bias parameter. It is easy to see that if $\beta \to 1$, then the consumption stream generated by this model is arbitrarily close to the ED-rationalizable behavior.

Definition 5. (Approximate s/m-rationalizability) We say that $(\boldsymbol{\rho}_t, \mathbf{c}_t^*)_{t\in\mathcal{T}}$ is **approximately consistent** with s/m-rationalizability if there exists a sequence of random variables $(\boldsymbol{\lambda}_k', \mathbf{v}_k')' \in \Lambda \times \mathbb{R}_+^{|\mathcal{T}|}$, $k = 1, 2, \ldots$, such that

$$\mathbb{P}\left(\mathbb{1}\left(\, \mathbf{v}_{k,t} - \mathbf{v}_{k,s} \geq \boldsymbol{\lambda}_{k,t}\boldsymbol{\rho}_t'[\mathbf{c}_t^* - \mathbf{c}_s^*]\,\right)\right) \to_{k\to+\infty} 1,$$

for all $s, t \in \mathcal{T}$.

4.2. Introducing Measurement Error

Theorem 1 and Lemma 2 allow us to get testable implications of s/m-rationalizability. These implications depend solely on the distribution of $\boldsymbol{\lambda} = (\boldsymbol{\lambda}_t)_{t\in\mathcal{T}}$ and $\mathbf{v} = (\mathbf{v}_t)_{t\in\mathcal{T}}$. The usual approach to testing s/m-rationalizability would amount to solving a linear programming problem corresponding to Theorem 1 at the level of individual consumers. However, this common practice does not work any more in the presence of **measurement error**. When **true consumption** is measured erroneously, we observe not $\mathbf{c}_t^*$ but rather a perturbed version of it.

Remark 3. Measurement error in consumption in (recall) surveys arises because consumers are asked to report past consumption patterns in a periodic nature. Consumers may fail to correctly report their consumption level due to bounded recall, social desirability, and mistakes in communication. Thus, self-reported consumption is likely to be mismeasured.

Define the **measurement error** $\mathbf{w} = (\mathbf{w}_t)_{t\in\mathcal{T}} \in W$, where $\mathbf{w}_t \in W \subseteq \mathbb{R}^L$, and $t \in \mathcal{T}$, as the difference between reported consumption, $\mathbf{c} = (\mathbf{c}_t)_{t\in\mathcal{T}}$, and true consumption, $(\mathbf{c}^*_t)_{t\in\mathcal{T}}$. That is,

$$\mathbf{w}_t = \mathbf{c}_t - \mathbf{c}^*_t,\ t \in \mathcal{T}. \tag{1}$$

It is important to note that we *define* the measurement error on consumption. We do not make any assumptions about how the difference between $\mathbf{c}$ and $\mathbf{c}^*$ arises (i.e., we allow for measurement error in consumption to be multiplicative or additive).[13]

By Lemma 2 we can immediately conclude that the observed $\mathbf{x} = (\boldsymbol{\rho}_t, \mathbf{c}_t)_{t\in\mathcal{T}}$ can be s/m-rationalized if and only if there exist $(\boldsymbol{\lambda}_t, \mathbf{v}_t, \mathbf{w}_t)_{t\in\mathcal{T}}$, with $(\boldsymbol{\lambda})_{t\in\mathcal{T}}$ supported on or inside Λ, such that

$$\mathbf{v}_t - \mathbf{v}_s \geq \boldsymbol{\lambda}_t \boldsymbol{\rho}'_t(\mathbf{c}_t - \mathbf{c}_s + \mathbf{w}_s - \mathbf{w}_t) \quad \text{a.s.}, \forall s, t \in \mathcal{T}.$$

However, we know that without restrictions on the distribution of measurement error, revealed-preference tests have no power. That is, there always exists a measurement error $\mathbf{w}$ such that the observed $\mathbf{x}$ is consistent with s/m-rationalizability. We therefore impose reasonably weak centering restrictions on measurement error, namely, mean budget neutrality. In this case, we show that the restrictions on measurement error interact with the revealed-preference restrictions in order to provide a test for s/m-rationalizability with asymptotic power equal to 1.

Assumption 1. *(Mean Budget Neutrality)* $\mathbb{E}\left[\boldsymbol{\lambda}_t \boldsymbol{\rho}'_t \mathbf{w}_t\right] = 0$, *for all* $t \in \mathcal{T}$.

Assumption 1 implies that measurement error does not alter the weighted mean value of total expenditure, $\mathbb{E}\left[\boldsymbol{\lambda}_t \boldsymbol{\rho}'_t \mathbf{c}_t\right] = \mathbb{E}\left[\boldsymbol{\lambda}_t \boldsymbol{\rho}'_t \mathbf{c}^*_t\right]$. The weights are interpreted as the marginal utility of income. In other words, mean budget neutrality captures the idea that consumers, on average, may remember the total expenditure level better than the actual details.

Our assumption about measurement error does not imply the classical measurement-error assumptions that may fail in the consumer environment. In fact, it is compatible with nonclassical measurement error, such as $\mathbb{E}\left[\mathbf{w}_t | \mathbf{c}^*_t\right] \leq 0$ (Carroll et al. (2014)). Assumption 1 is implied by the following conditions: (i) $\mathbb{E}\left[\, \boldsymbol{\rho}'_t \mathbf{w}_t | \mathbf{c}^*_t \,\right] = 0$, and (ii) $(\boldsymbol{\lambda}_t \perp\!\!\!\perp \mathbf{w}_t) | \mathbf{c}^*_t$. Condition (i) means that the consumers correctly recall, on average, their total expenditure on all goods conditional on true consumption, and condition (ii) means that measurement error is independent of the random marginal utility of wealth, conditional on true consumption.[14]

[13]Formally this makes the support W depend on the support of both the observed and the true consumption. For simplicity we omit this dependency from the notation.

[14]Mean budget neutrality will fail if measurement error in consumption is systematic (i.e., every consumer over-reports expenditure or under-reports expenditure simultaneously). There is some evidence for the case of recall surveys, such as the one we use in our application, that measurement error in consumption can be over-reported or under-reported in a non-systematic fashion (Mathiowetz et al. (2002)). We are not interested in systematic

Example 2. (Additive Measurement Error) Consider a case where $\mathbf{c}_{t,l} = \mathbf{c}^*_{t,l} + \boldsymbol{\epsilon}_{t,l}$ for $l = 1, \ldots, L$; and where $\boldsymbol{\epsilon}_{t,l} \sim TN_{[-a,a]}(0, \sigma)$ for $l = 1, \ldots, L-1$ (from a truncated normal with variance σ and bounds $[-a, a]$ for some positive $a > 0$) such that $\mathbf{c}_{t,l} \geq 0$ a.s.. Assume, moreover, that $\boldsymbol{\epsilon}_{t,L} = -\frac{1}{\rho_{t,L}} \sum_{l=1}^{L-1} \boldsymbol{\rho}_{t,l} \boldsymbol{\epsilon}_{t,l}$ a.s.. In words, the consumers report correctly their total expenditure with probability 1 while making some mistakes about the budget shares for each good. Measurement error is defined as $\boldsymbol{w}_{t,l} = \boldsymbol{\epsilon}_{t,l}$. Note that $\mathbf{w}_{t,l}$ is independent from $\boldsymbol{\lambda}_t$ conditional on $\mathbf{c}^*_t$, and by construction $\sum_{l=1}^{L} \boldsymbol{\rho}_{t,l} \boldsymbol{\epsilon}_{t,l} = 0$ a.s.. This means that Assumption 1 holds.

Example 3. (Price and Expenditures with Measurement Error) Our definition of measurement error is general enough to cover special cases of measurement error in prices and expenditures. Say that both reported good-level expenditures $\boldsymbol{\rho}_{t,l}\mathbf{c}_{t,l}$, and transformed prices $\boldsymbol{\rho}_{t,l}$ are measured with error. In particular, the error enters multiplicatively, such that $\boldsymbol{\rho}_{t,l}\mathbf{c}_{t,l} = \boldsymbol{\rho}^*_{t,l}\mathbf{c}^*_{t,l}\boldsymbol{\epsilon}_{1,t,l}$, and $\boldsymbol{\rho}_{t,l} = \boldsymbol{\rho}^*_{t,l}\boldsymbol{\epsilon}_{2,l}$ (with $\boldsymbol{\epsilon}_i$ for $i = 1, 2$ denoting positive random vectors). In practice, observed consumption is computed as the ratio of these mismeasured expenditures and prices: $\mathbf{c}_{t,l} = \boldsymbol{\rho}^*_{t.l}\mathbf{c}^*_{t,l}\boldsymbol{\epsilon}_{1,t,l}/\boldsymbol{\rho}^*_{t,l}\boldsymbol{\epsilon}_{2,l}$. The result is that observed consumption is also mismeasured and is of the form $\mathbf{c}_{t,l} = \mathbf{c}^*_{t,l}\boldsymbol{\epsilon}_{3,t,l}$, where $\boldsymbol{\epsilon}_{3,t,l} = \boldsymbol{\epsilon}_{1,t,l}/\boldsymbol{\epsilon}_{2,l}$. Here we assume that the price measurement error has to be time-invariant. Note that by defining measurement error as $\hat{\mathbf{w}}_{t,l} = \mathbf{c}^*_{t,l}\boldsymbol{\epsilon}_{3,t,l} - \mathbf{c}^*_{t,l}/\boldsymbol{\epsilon}_{2,l}$ instead of $\mathbf{w}_{t,l} = \mathbf{c}^*_{t,l}\boldsymbol{\epsilon}_{3,t,l} - \mathbf{c}^*_{t,l}$, we can rationalize the data set $(\boldsymbol{\rho}_t, \mathbf{c}_t)_{t\in\mathcal{T}}$, because the equality $\boldsymbol{\rho}'_t(\mathbf{c}_t - \mathbf{c}_s + \hat{\mathbf{w}}_s - \hat{\mathbf{w}}_t) = \sum_{l=1}^{L} \boldsymbol{\rho}^*_{t,l}\boldsymbol{\epsilon}_{2,l}(\mathbf{c}^*_{t,l}/\boldsymbol{\epsilon}_{2,l} - \mathbf{c}^*_{s,l}/\boldsymbol{\epsilon}_{2,l})$ simplifies to $\boldsymbol{\rho}'_t(\mathbf{c}_t - \mathbf{c}_s + \hat{\mathbf{w}}_s - \hat{\mathbf{w}}_t) = (\boldsymbol{\rho}^*_t)'(\mathbf{c}^*_t - \mathbf{c}^*_s)$. Our methodology works without any change for this case since $\mathbf{w}_t$ and $\hat{\mathbf{w}}_t$ are latent random variables from the econometrician's point of view, and their only role is to perturb the data set to make it consistent with the revealed-preference restrictions. This means that the case presented in this example with price measurement error is indistinguishable from a case with a perfectly measured price vector and error in measuring consumption when measurement error in prices is time-invariant. Assumption 1 can be satisfied under standard conditional independence assumptions between $\boldsymbol{\lambda}_t$, $\boldsymbol{\epsilon}_{1,t,l}$, and $\boldsymbol{\epsilon}_{2,l}$, in addition to a restriction that the conditional mean of all multiplicative disturbance terms is equal to 1. Evidently, if a random array $\mathbf{x} = (\boldsymbol{\rho}_t, \mathbf{c})_{t\in\mathcal{T}}$ is s/m-rationalizable without measurement error in prices, it will be rationalizable, allowing for measurement error in prices.

In the appendix, we provide a variant of our methodology to deal with measurement error both in prices and consumption.

measurement error as it may be confounded with behavioral mistakes. For instance, time-inconsistency may cause over-consumption in the present, which may be confounded with systematic measurement error under the null hypothesis of exponential discounting behavior).

5. Econometric Framework

5.1. Characterization of the Model via Moment Conditions

Now, we recast the empirical content of the revealed-preference inequalities in a form amenable to statistical testing. In particular, we write a set of moment conditions that will summarize the empirical content of s/m-rationalizability. Recall that $\mathbf{x} \in X$ denotes observed quantities. Let $\mathbf{e} = (\boldsymbol{\lambda}', \mathbf{v}', \mathbf{w}')' \in E$ denote the vector of latent random variables. The support E depends on the fixed support Λ that characterizes the particular model of interest. We use $\mathcal{P}_X$, $\mathcal{P}_{E,X}$, and $\mathcal{P}_{E|X}$ to denote the set of all probability measures defined over the support of $\mathbf{x}$, $(\mathbf{e}', \mathbf{x}')'$, and $\mathbf{e}|\mathbf{x}$, respectively. (Recall that the boldface font letters denote random objects.) Define the following moment functions:

$$\begin{aligned}
g_{I,t,s}(\mathbf{x}, \mathbf{e}) &= \mathbb{1}\left(\mathbf{v}_t - \mathbf{v}_s - \boldsymbol{\lambda}_t \boldsymbol{\rho}_t'[\mathbf{c}_t - \mathbf{w}_t - \mathbf{c}_s + \mathbf{w}_s] \geq 0\right) - 1, \quad t \neq s \in \mathcal{T};\\
g_{B,t}(\mathbf{x}, \mathbf{e}) &= \boldsymbol{\lambda}_t \boldsymbol{\rho}_t' \mathbf{w}_t, \quad t \in \mathcal{T};\\
g_I(\mathbf{x}, \mathbf{e}) &= (g_{I,t,s}(\mathbf{x}, \mathbf{e}))_{t \neq s \in \mathcal{T}};\\
g_B(\mathbf{x}, \mathbf{e}) &= (g_{B,t}(\mathbf{x}, \mathbf{e}))_{t \in \mathcal{T}};\\
g(\mathbf{x}, \mathbf{e}) &= (g_I(\mathbf{x}, \mathbf{e})', g_B(\mathbf{x}, \mathbf{e})')'.
\end{aligned}$$

We end up having $k = |\mathcal{T}|^2 - |\mathcal{T}|$ and $q = |\mathcal{T}|$ moment functions which correspond to inequality conditions (g_I) and the mean budget-neutrality conditions (g_B), respectively. Define $\mathbb{E}_{\mu \times \pi}\left[\, g(\mathbf{x}, \mathbf{e})\,\right] = \int_X \int_{E|X} g(x, e) d\mu d\pi$, where $\mu \in \mathcal{P}_{E|X}$ and $\pi \in \mathcal{P}_X$.

Theorem 2. *The following are equivalent:*

(i) A random vector $\mathbf{x} = (\boldsymbol{\rho}_t, \mathbf{c}_t)_{t \in \mathcal{T}}$ is approximately s/m*-rationalizable such that Assumption 1 holds.*

(ii)

$$\inf_{\mu \in \mathcal{P}_{E|X}} \left\| \mathbb{E}_{\mu \times \pi_0}\left[\, g(\mathbf{x}, \mathbf{e})\,\right] \right\| = 0,$$

where $\pi_0 \in \mathcal{P}_X$ is the observed distribution of $\mathbf{x}$.

Theorem 2 establishes the equivalence between (i) s/m-rationalizability with the mean budget-neutral measurement error and (ii) a system of moment conditions. In other words, the observed consumption pattern, captured by the random array $(\boldsymbol{\rho}_t, \mathbf{c}_t)_{t \in \mathcal{T}}$, can be s/m-rationalized under the restrictions on measurement error if and only if there exists a distribution of latent variables conditional on observables that satisfies the revealed-preference inequalities with probability 1, for the given support Λ.

Our statistical notion of rationalizability makes clear that when one is dealing with measurement error, no revealed-preference test can decide whether a finite sample is consistent with model m. We can decide only that the data set is asymptotically consistent with the model as the sample size goes to infinity. Moreover, even asymptotically, there is no way to differentiate between the notion of approximate s/m-rationalizability and the notion of exact s/m-rationalizability. Nonetheless, we can do traditional hypothesis testing and decide at a fixed confidence level whether we reject the null hypothesis of (approximate) model m consistency under Assumption 1 for a given sample. Conceptually, our notion of rationalizability corresponds to the extended notion of an identified set in Schennach (2014).

Note that the test is not yet formally established. We have a set of latent random variables $\mathbf{e}$ distributed according to an unknown $\mu \in \mathcal{P}_{E|X}$. This problem can be solved nonparametrically using the the Entropic Latent Variable Integration via Simulation (ELVIS) of Schennach (2014). The main advantage of the ELVIS approach is that it allows us to formulate a test that can be implemented in survey data suffering from measurement error of the type described only in terms of observables.

5.2. ELVIS and Its Implications for Testing and Inference

We start this section by showing how the nonparametric results of Theorem 2 can be used to construct a set of (equivalent) parametric maximum-entropy moment conditions (MEM conditions) using Schennach (2014). Next, we provide a semi-analytic solution to the MEM conditions. Finally, we propose a procedure to test for s/m-rationalizability. Following Schennach (2014), we define the Maximum-Entropy Moment (MEM) as follows:

Definition 6. (Maximum-Entropy Moment, MEM) The MEM of the moment $g(x, e)$, for a fixed x, is

$$h(x;\gamma) = \frac{\int_{e\in E|X} g(x,e)\exp(\gamma' g(x,e))d\eta(e|x)}{\int_{e\in E|X}\exp(\gamma' g(x,e))d\eta(e|x)},$$

where $\gamma \in \mathbb{R}^{k+q}$ is a nuisance parameter, and $\eta(\cdot|\cdot) \in \mathcal{P}_{E|X}$ is an arbitrary user-input distribution function supported on $E|X$ such that $\mathbb{E}_{\pi_0}\left[\log \mathbb{E}_{\eta(\cdot|\cdot)}\left[\exp(\gamma' g(\mathbf{x},\mathbf{e}|\mathbf{x}))\right]\right]$ exists and is twice continuously differentiable in γ for all $\gamma \in \mathbb{R}^{k+q}$.

In words, the MEM h is the marginal moment of the function g, at which the latent variable has been integrated out. The MEM depends only on the observable random variables.[15] The importance of the MEM is captured in the following result.

Theorem 3. *The following are equivalent:*

(i) A random array $\mathbf{x} = (\boldsymbol{\rho}_t, \mathbf{c}_t)_{t\in\mathcal{T}}$ is approximately s/m*-rationalizable such that Assumption 1 holds.*

(ii)

$$\inf_{\gamma\in\mathbb{R}^{k+q}} \|\mathbb{E}_{\pi_0}\left[h(\mathbf{x};\gamma)\right]\| = 0,$$

where $\pi_0 \in \boldsymbol{P}_X$ is the observed distribution of $\mathbf{x}$.

[15]The choice of the $\eta(e|x)$ is inessential since it affects only the nuisance parameter value (Schennach (2014)).

We emphasize that Theorem 3 provides both *necessary and sufficient* conditions for the observed data to be (approximately) s/m-rationalizable. This represents an important gain in power with respect to any of the averaging-based tests of revealed-preference models that are usually used in the presence of measurement error.

We must stress that we could impose high-level technical assumptions to ensure that the sequence of random latent variables that approximates model m converges to a proper random variable. Thus, this limiting random variable would ensure (i) that the infimum in Theorem 3 is attained, and (ii) that the notion of approximate rationalizability collapses to exact rationalizability. However, this obscures the fact that any assumption made in that direction has no testable implications.

The remarkable advantage of applying the results of Schennach (2014) to the revealed-preference approach is that it marginalizes out the latent random variables. More importantly, we have a robust statistical framework with which to test our models in the presence of measurement error. In particular, we have not made any strong distributional assumptions about $\boldsymbol{\lambda}_t$ or $\mathbf{u}$ (the heterogeneous tastes). The only assumptions are the concavity of the utility function, and a centering assumption about measurement error. In short, the proposed methodology allows us to test s/m-rationalizability in a robust manner without parametric assumptions about preferences or strong distributional assumptions about measurement error.[16]

Remark 4. Theorem 3 does not imply that the distribution of the latent variables (or their support) is point-identified. In fact, it will always be set-identified.

5.3. Semi-analytic Solution for the MEM for s/m-Rationalizability with Measurement Error

One can directly employ the MEM in Theorem 3 to test model m. However, doing so is potentially problematic. One possible concern is the fact that the number of MEM, $k + q = |\mathcal{T}|^2$, grows quadratically with $|\mathcal{T}|$. Moreover, γ_0, the nuisance parameter value at which infimum is achieved, may be set-identified when unbounded (e.g., some of the components of γ_0 may be equal to infinity), which would therefore lead to nonstandard testing procedures.

Here we show that there exists a semi-analytic solution to the optimization problem where every component of γ_0 that corresponds to the revealed-preference inequality constraints is equal to $+\infty$, and every component of γ_0 that corresponds to the mean budget-neutrality constraint is finite and unique. Thus, for testing purposes (under the null hypothesis of model m), we can restrict the first k components of γ to be equal to $+\infty$, and then minimize a convex objective function over a parameter space of lower dimensionality.

Assumption 2. *(Nondegeneracy) The distribution of $\boldsymbol{\rho}_t'\mathbf{w}_t$ is nondegenerate for all $t \in \mathcal{T}$.*

Assumption 3. *(Bounded support) The random array $\mathbf{x} = (\boldsymbol{\rho}_t, \mathbf{c}_t)_{t\in\mathcal{T}}$ has a bounded support.*

Assumption 2 rules out cases in which there is no measurement error. Our methodology still works for cases without measurement error, but in those cases, it is preferable to use the equivalent

[16] At this point, we can use as an alternative the methodology presented by Ekeland et al. (2010) to deal with latent variables in our moment conditions.

deterministic revealed-preference benchmark. Assumption 3 is made to simplify the analysis and can be replaced by tail restrictions on the distribution of $\mathbf{x}$.

Note that the user-specified distribution $\eta(\cdot|\cdot)$ should obey the same restrictions as the unknown distribution of latent $\mathbf{e}$. Thus, we impose the following restrictions on $\eta(\cdot|\cdot)$:

Definition 7. (User-specified MEM distribution) Almost surely in $\mathbf{x}$, the user-specified distribution $\eta(\cdot|\mathbf{x})$ satisfies all of the following:

(i) The set $\tilde{E}|X = \{e \in E|X \ : \ g_I(\mathbf{x}, e) = 0\}$ has a positive measure under $\eta(\cdot|\mathbf{x})$.

(ii) There exist two subsets of $\tilde{E}|X$, E' and E'', with a positive measure under $\eta(\cdot|\mathbf{x})$, such that componentwise $\sup_{e\in E'} g_B(\mathbf{x}, e) < 0 < \inf_{e\in E''} g_B(\mathbf{x}, e)$.

(iii) For every finite $\gamma_B \in \mathbb{R}^q$,

$$\int_{E|X} \|g_B(\mathbf{x}, e)\|^2 \exp(\gamma_B' g_B(\mathbf{x}, e)) d\eta(e|\mathbf{x}) < \infty.$$

The first condition in Definition 7 requires that the support of η allows the inequalities to be satisfied. The second and third conditions are regularity conditions. This is a definition and not an assumption, as we can always construct such an $\eta(\cdot|\cdot)$.[17] We are ready to present our main result.

Theorem 4. *Given a user-specified measure $\eta(\cdot|\cdot)$ that satisfies the three conditions in Definition 7, the following are equivalent:*

(i) A random array $\mathbf{x} = (\boldsymbol{\rho}_t, \mathbf{c}_t)_{t\in\mathcal{T}}$ is approximately consistent with s/m*-rationalizability such that assumptions 1-3 hold.*

(ii) For any sequence $\{\gamma_{I,l}\}_{l=1}^{+\infty}$ that componentwise diverges to $+\infty$,

$$\lim_{l\to+\infty} \min_{\gamma_B\in\mathbb{R}^q} \left\| \mathbb{E}_{\pi_0} \left[h(\mathbf{x}; (\gamma_{I,l}', \gamma_B')') \right] \right\| = 0. \tag{2}$$

The sequence of minimizers of (2), $\{\gamma_{B,l}\}$, converges to some finite $\gamma_{0,B}$ that does not depend on $\{\gamma_{I,l}\}_{l=1}^{+\infty}$.

(iii)

$$\min_{\gamma_B\in\mathbb{R}^q} \left\| \mathbb{E}_{\pi_0} \left[\tilde{h}_B(\mathbf{x}; \gamma_B) \right] \right\| = 0, \tag{3}$$

where

$$\tilde{h}_B(x; \gamma) = \frac{\int_{e\in E|X} g_B(x, e) \exp(\gamma' g_B(x, e)) \mathbb{1}\left(g_I(x, e) = 0 \right) d\eta(e|x)}{\int_{e\in E|X} \exp(\gamma' g_B(x, e)) \mathbb{1}\left(g_I(x, e) = 0 \right) d\eta(e|x)}.$$

Moreover, the minimizer of (3) is finite, and is equal to $\gamma_{0,B}$.

Theorem 4 substantially simplifies the conclusion of Theorem 3. First, we need to minimize the globally convex objective function over a much smaller parameter space ($\mathbb{R}^q = \mathbb{R}^{|\mathcal{T}|}$ instead

[17] Schennach (2014) provides a generic construction that can be used here.

of $\mathbb{R}^{k+q} = \mathbb{R}^{|\mathcal{T}|^2}$). Thus, the problem becomes computationally tractable. Second, if the data is consistent with s/m-rationalizability, then the minimizer, $\gamma_{0,B}$, has to be *finite and unique*. Finally, for implementation purposes, it suffices to replace γ_I by a very large number.[18] The intuition behind Theorem 4 is that the revealed-preference inequalities presented here restrict only the conditional support of the latent variables (including the measurement error in consumption). Hence, given the support restrictions captured by the revealed-preference inequalities, only the centering condition comes in the form of moments.

5.4. Testing

Theorem 4 provides moment conditions that are necessary and sufficient for the data $\{\mathbf{x}_i\}_{i=1}^n = \{(\boldsymbol{\rho}_{t,i}, \mathbf{c}_{t,i})_{t\in\mathcal{T}}\}_{i=1}^n$ (where n is the sample size), to be approximately consistent with s/m-rationalizability. Now, define the following sample analogues of the MEM and the MEM-variance matrix:

$$\hat{\tilde{h}}_B(\gamma) = \frac{1}{n}\sum_{i=1}^n \tilde{h}_B(\mathbf{x}_i, \gamma);$$
$$\hat{\tilde{\Omega}}(\gamma) = \frac{1}{n}\sum_{i=1}^n \tilde{h}_B(\mathbf{x}_i, \gamma)\tilde{h}_B(\mathbf{x}_i, \gamma)' - \hat{\tilde{h}}_B(\gamma)\hat{\tilde{h}}_B(\gamma)'.$$

Let Ω^- denote the generalized inverse of the matrix Ω. The testing procedure we propose is due to Schennach (2014), and is based on this test statistic:

$$\mathrm{TS}_n = n \inf_{\gamma\in\mathbb{R}^q} \hat{\tilde{h}}_B(\gamma)'\hat{\tilde{\Omega}}^-(\gamma)\hat{\tilde{h}}_B(\gamma).$$

Assumption 4. *The data $\{\mathbf{x}_i\}_{i=1}^n$ is i.i.d.*

Theorem 5. *Suppose assumptions 1–4 hold. Then under the null hypothesis that the data is approximately consistent with* s/m-*rationalizability, it follows that*

$$\lim_{n\to\infty} \mathbb{P}\left(\mathrm{TS}_n > \chi^2_{q,1-\alpha}\right) \leq \alpha, \tag{4}$$

for every $\alpha \in (0,1)$.

If, moreover, the minimal eigenvalue of the variance matrix $\mathbb{V}[\tilde{h}_B(\mathbf{x},\gamma)]$ is uniformly, in γ, bounded away from zero and the maximal eigenvalue of $\mathbb{V}[\tilde{h}_B(\mathbf{x},\gamma)]$ is uniformly, in γ, bounded from above, then, under the alternative hypothesis that the data is not approximately consistent with s/m-*rationalizability, it follows that*

$$\lim_{n\to\infty} \mathbb{P}\left(\mathrm{TS}_n > \chi^2_{q,1-\alpha}\right) = 1.$$

[18] Alternatively, we can generate the new measure $d\tilde{\eta}(\cdot|x) = \mathbb{1}\left(g_I(x,\cdot) = 0\right) d\eta(\cdot|x)$ by sampling from $\eta(\cdot|x)$ and then accepting a draw only if the revealed-preference inequalities captured by $\mathbb{1}\left(g_I(x,\cdot) = 0\right)$ are satisfied. The last part usually amounts to solving a linear program. For details about the implementation see section the appendix.

6. Empirical Application: Testing Exponential Discounting in Survey Data

In this section, we use our methodology to test the exponential discounting hypothesis in a consumer panel data set gathered from single-individual and couples' households in Spain. In particular, we work with the data set used in Adams et al. (2014): the Spanish Continuous Family Expenditure Survey (*Encuesta Continua de Presupuestos Familiares* 1985-1997). The data set consists of the expenditures for 185 individuals and 2004 couples, as well as prices for 17 commodities (categories of goods) recorded over four consecutive quarters. The categories of goods are: all food and nonalcoholic drinks, all clothing, cleaning, nondurable articles, household services, domestic services, public transport, long-distance travel, other transport, petrol, leisure (four categories), other services (two categories), and food consumed outside the home. The data set also contains information on the nominal interest rate on consumer loans faced by the household in any particular quarter.[19]

Some notable studies in the deterministic revealed-preference context that have used this Spanish household consumer survey are Beatty and Crawford (2011), Blow et al. (2013) and Adams et al. (2014). The conclusion that we drawn from applying the deterministic methodology used in Browning (1989) and Blow et al. (2017) to the sample of interest is that the exponential discounting hypothesis is rejected for a substantial fraction of single-individual and couples' households. That is, a sizable fraction of households behave in a manner inconsistent with the predictions of the model. In contrast, we find, at the 95 percent confidence level, that the exponential discounting model with measurement error cannot be rejected for single-individual households. This fact indicates that deterministic tests may not be very informative about the behavior in a population, due to measurement error. Small violations of the deterministic revealed-preference inequalities will lead to big rejection rates. Introducing measurement error into the analysis takes these small violations into account.[20]

To the best of our knowledge, we are the first to perform traditional statistical hypothesis testing of the exponential discounting consumer model in the presence of measurement error in a completely nonparametric fashion. Our empirical findings suggest that practitioners of the revealed-preference methodology should take into account measurement error in order to be able to conduct traditional hypothesis testing. Furthermore, we wish to emphasize that our methodology does not require identification of the distribution of the discount factor. Nevertheless, we still can construct confidence sets for its support.

Formally, we test for s/ED-rationalizability with (i) effective prices equal to the discounted spot prices, $\boldsymbol{\rho}_t = \boldsymbol{\rho}_t^{\mathrm{ED}}$ (defined in Table 1), and (ii) random marginal utility of income equal to the discounted value of one unit of wealth, $\boldsymbol{\lambda}_t = \mathbf{d}^{-t}$, where $\mathbf{d}$ is interpreted as the random discount factor supported on or inside $(0, 1]$.

We allow for flexible support of the random discount factor, and we assume that it is compact.

[19] We spare the reader more details and refer them instead to Adams et al. (2014) for further information on the data set.

[20] In the appendix, we also establish that our test fails to rejects a special version of the collective household consumption problem presented in Adams et al. (2014) (see section 12).

Formally, denote $\mathbf{d}$ as the random discount factor that depends on parameter θ_0. We impose Assumption 5.

Assumption 5. *(Compact Support) The random discount factor has support contained in* $D_{\theta_0} = [\theta_0, 1]$*, where* $0 < \epsilon \leq \theta_0 \leq 1$ *for some* $\epsilon > 0$*.*

The parametrization in Assumption 5 is not restrictive, and can easily be replaced by any parametric restriction on the support of the discount factor by allowing a flexible upper bound of the support. For instance, instead of a compact interval, one could assume that support is discrete or a finite union of disjoint compact intervals. There is some evidence in experimental setups that supports the restriction that the upper bound of the discount factor distribution is close to unity (e.g., Montiel Olea and Strzalecki (2014)); thus, we choose the upper bound to be equal to 1. This is a weak assumption that helps to facilitate the visualization of the present method and simplifies the optimization procedure. More important, our results are still valid even if this assumption does not hold. In case this assumption fails and the null hypothesis is not rejected, we can only conclude that there is a distribution of the discount factor with support on or inside $[\theta_0, 1]$. (Note that θ_0 is not a deterministic discount factor, but rather the lower bound of the support of the random discount factor $\mathbf{d}$.)

We also require that the random array $\mathbf{x} = (\boldsymbol{\rho}_t, \mathbf{c}_t)_{t\in\mathcal{T}}$ has bounded support. For a fixed θ_0, we can obtain the support Λ of the random vector $(\boldsymbol{\lambda}_t)_{t\in\mathcal{T}}$ with $\boldsymbol{\lambda}_t = \boldsymbol{d}^{-t}$. We can test for s/ED($\theta_0$)-rationalizability, where the dependence on the lower bound on the support of the marginal utility of income is made explicit. Note that under Assumption 5 for given $\epsilon > 0$, the support of the marginal utility of income $\boldsymbol{\lambda}_t$ is bounded above by ϵ^{-t} (and bounded below by 1).[21] Imposing the additional Assumptions 1, 2, and 4, we can apply our testing methodology developed in section 5.4. Assumption 1 implies that measurement error does not alter the mean discounted value of total expenditure, $\mathbb{E}\left[\mathbf{d}^{-t}\boldsymbol{\rho}_t'\mathbf{c}_t\right] = \mathbb{E}\left[\mathbf{d}^{-t}\boldsymbol{\rho}_t'\mathbf{c}_t^*\right]$.

6.1. The Results

Single-Individual Households

We apply the deterministic methodology of Browning (1989) to single-individual households. Our initial conclusion is that 81.1 percent of the single-individual households behave inconsistently with exponential discounting (even when allowing for substantially more heterogeneity than previous works).[22] Next, we revisit this conclusion using our methodology, which addresses measurement error, while allowing a heterogeneous discount factor. We find that, at least at the 95 percent confidence level, we cannot reject exponential discounting. Formally, we find at least at the 95 percent confidence level that the random array $\mathbf{x} = (\boldsymbol{\rho}_t, \mathbf{c}_t)_{t\in\mathcal{T}}$ is s/ED-rationalizable with a random

[21] For this particular application we choose $\epsilon = 0.10$.

[22] We search for each individual household discount factor d in the grid $\{0.1, 0.15, \cdots, 1\}$. See Crawford (2010), Adams et al. (2014) and Blow et al. (2017) for discount factor ranges close to $[0.9, 1]$.

discount factor $\mathbf{d}$ supported on or inside $[0.1, 1]$. We are also interested in learning more information about the support of the random discount factor. For that reason, for every fixed θ_0 on the grid $\{0.1, 0.15, \ldots, 1\}$, we compute $\mathrm{TS}_n(\theta_0)$ and compare it with $\chi^2_{4,0.95} = 9.5$.[23] The results are presented in figure 1. The smallest value of the lower bound of the support of the discount factor to pass the test is 0.1 ($\mathrm{TS}_n(0.1) = 3.05$); for $\theta_0 \geq 0.55$ the test is rejected.[24] That means that the support of $\mathbf{d}$ cannot be a subset of $[0.55, 1]$, with at least a 95 percent confidence level.[25]

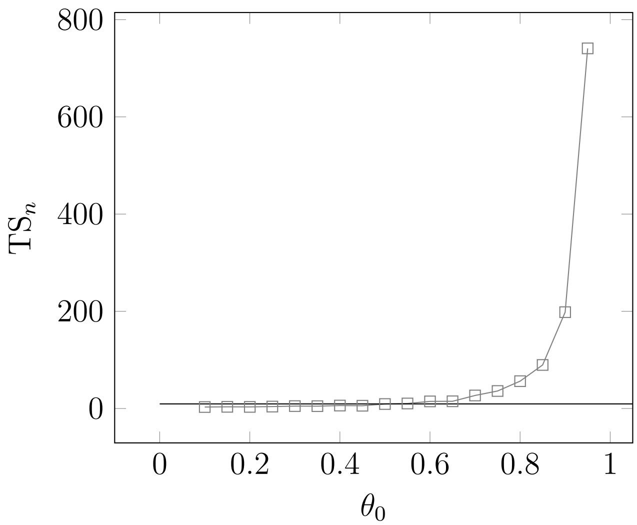

Figure 1 – Values of TS_n. The horizontal line corresponds to the 95 percent quantile of χ^2_4.

Couples' Households

For the couples' households, the deterministic test of Browning (1989) rejects the exponential discounting model for 88.5 percent of the observations. Although this number seems large, one should keep in mind that for single-individual households the same deterministic test rejects the model in 81.1 percent of the cases. At the same time, our method does not reject the exponential discounting model for single-individual households. But we do reject the model for couples' households. In the case of couples' households, the test statistic is above the conservative critical value based on the χ^2_4 distribution.

The test statistic takes its minimum value in the grid search at $\theta_0 = 0.1$.[26] The value of the test statistic is $\mathrm{TS}_n(0.1) = 9.86$ (the p-value is 0.043). This is above the conservative critical value

[23] We again make explicit the dependence of the test statistic on θ_0.

[24] The p-value monotonically decreases from 0.548 for $\theta_0 = 0.1$ to 0 for $\theta_0 = 0.95$. The p-values for $\theta_0 = 0.5$ and $\theta_0 = 0.55$ are 0.056 and 0.035, respectively.

[25] In other words, we reject the null hypothesis that the random array $\mathbf{x} = (\boldsymbol{\rho}_t, \mathbf{c}_t)_{t \in \mathcal{T}}$ is s/ED-rationalizable with a random discount factor supported on or inside $[0.55, 1]$. We provide a pseudo-algorithm of our testing procedure in the appendix.

[26] We searched for $\theta_0 \in [0.1, 1]$, with a grid $\{0.1, 0.15, 0.2, 0.4, 0.7, 0.9\}$. Since we reject for $\theta_0 = 0.1$, we do not need a finer grid, since exponential discounting will be rejected for the rest of the points, as the numerical exercise confirms.

$\chi^2_{4,0.95} = 9.5$.[27] Hence, we reject the null hypothesis of exponential discounting for all values of θ_0 at the 95 percent confidence level.

6.2. Discussion and Related Work on Testing the Exponential Discounting Model

In our empirical application, we found support for exponential discounting behavior for single-individual households; this was in contrast with the result of applying Browning (1989) methodology to our sample. At the same time, we reject exponential discounting behavior for the case of couples' households. This supports the negative finding in Blow et al. (2017) for a very similar sample using a deterministic approach. The support we found for exponential discounting behavior for single-individual households may be surprising, given the experimental evidence against exponential discounting for experimental subjects. However, we believe that this result makes sense, given that the decision task that we study (i.e., periodically allocating a budget share to a good category for 4 consecutive time periods) for the following reasons. First, the decision task involves higher stakes than an experimental decision task (Halevy (2015)). Second, the consumer may have developed some expertise at the decision task given its repetitive nature (Choi et al. (2014)). Third, the consumer may be time-consistent for budgeting decisions such as how much of the budget share will be allocated to food versus recreation every month, yet may be time-inconsistent in more granular decisions.[28]. We can think of a two-system decision-making framework with a mental-budgeting first step that is used to decide budget shares for big categories of goods, and a second (possibly) time-inconsistent system that uses the first state budgeting as a given constraint (Sulka (2017)).

One possible concern about our methodology is that its power is low, given its nonparametric nature. However, in the appendix (Monte Carlo) we report, for 1000 trials with a sample size of 1500, a rejection rate greater than or equal to 69 percent (with a data-generating process consistent with hyperbolic discounting). At the same time, when the data-generating process is consistent with exponential discounting, as expected, the rejection rate is close to the theoretical 5 percent (7.9 percent).[29]

Our results for couples' households provide the first nonparametric evidence which is robust to measurement error and which demonstrates that not all couples' households manifests behavior consistent with exponential discounting. We must emphasize that s/ED-rationalizability is rejected with a very weak assumption about measurement error in the couples' case, while it is not rejected in the single-individual household case. This should convince practitioners about the importance of modeling intrahousehold decision-making when dealing with intertemporal choice. While the exponential discounting model seems to be a good model for single individuals, it fails for couples.

[27] The test statistic for the rest of the grid jumps to 11.94 (the p-value is 0.018) for $\theta_0 = 0.15$, and then is increasing in the values of the grid $\{0.2, 0.4, 0.7, 0.9\}$ until it reaches 1084.279.

[28] For example the consumer may be more prone to temptation when deciding how much ice cream she wants to buy when she visits the supermarket, but may remain time-consistent in terms of the total expenditure for food in a given month

[29] In addition, our methodology rejects the null hypothesis of s/ED-rationalizability for single-individual households when the heterogeneity of the random discount factor is too small.

In the appendix, we establish that a suitable extension of our methodology, fails to reject the collective household consumption problem presented in Adams et al. (2014). In the model in Adams et al. (2014), each member of a couples' household is consistent with s/ED-rationalizability under these two assumptions: (i) full efficiency in the provision of household public goods, and (ii) common support for the distribution of preferences, and random discount factors (as well as relative household power, as captured by random Pareto weights to be defined in section 12). Given the support for exponential discounting for single-individual households, this finding is not surprising.[30]

The rejection of exponential discounting behavior for couples' households can be better understood given new theoretical results that show that aggregating time-consistent preferences may lead to time-inconsistent behavior (Jackson and Yariv (2015)).

The deterministic methodology of Browning (1989) concludes that 81.1 percent of single-individual households are inconsistent with the exponential discounting model, while, concluding that 88.5 percent of couples' households are inconsistent with this model. The fraction of households that is inconsistent with exponential discounting under the deterministic test is similar for both cases, but our statistical test rejects in the latter case while reaching the opposite conclusion in the former case. Consider a hypothetical case in which all households pass the deterministic test; in that case, they would also pass our stochastic test. Given this fact, the rejection of exponential discounting is driven by the fraction of couples' households that is not consistent with the deterministic test. The difference in conclusions is due to the fact that our test implicitly takes into account the severity of the violations of exponential discounting, and imposes the mean budget-neutrality assumption on the measurement error corrections. Clearly, for the case of single-individual households, the violations of the empirical implications of exponential discounting must be either small **or** compatible with the mean budget neutrality (i.e., the violations are non-systematic). This is not the case for couples' households. We emphasize that goodness-of-fit measures such as those used in Varian (1990) (e.g., Afriat's Cost Efficiency Index and the Money Pump) are not comparable to our approach. Our methodology will accept the consistency of a data set with a given model when violations are high according to goodness-of-fit measures, if those violations are non-systematic (i.e., the violations are compatible with mean budget neutrality).

Our findings for the support of the random exponential discount factor for the single-individual households imply that we cannot reject the fact that the lower bound of the heterogeneous discount factor θ_0 belongs to $[0.1, 0.55)$ (i.e., the random discount factor is supported on or inside $[\theta_0, 1]$). This lower bound is substantially smaller than the lower bounds previously considered in the literature (e.g., Adams et al. (2014), and Blow et al. (2017) consider a lower bound close to 0.9).[31] This lower bound is, however, consistent with experimental evidence that has found some experimental subjects with low discount factors (Montiel Olea and Strzalecki (2014),Echenique et al. (2014)). Our methodology does not permit the unique recovery of the distribution of the random discount factor.[32] In this sense, having low discount factors in the support does not imply that the fraction

[30]We must emphasize that the test for the collective model uses only sufficient conditions, as opposed to our test for exponential discounting that uses necessary and sufficient conditions. We leave as an avenue for future research the formal study of collective household decision-making within the framework proposed by Adams et al. (2014).

[31]Given that the deterministic tests typically reject exponential discounting, this is not informative.

[32]However, we can still ask questions about some characteristics of the random discount factor distribution, such

of such households is sizable. We emphasize that allowing for a broader support for the discount factor weakens the discriminatory power of the test, which makes the rejection of exponential discounting for couples a severe one. We conclude that both (large) heterogeneity of the random discount factors and measurement error are needed in order to s/ED-rationalize the sample of single-individual households.

Beatty and Crawford (2011) test for static utility maximization in a similar data set using a predictive success measure at the individual household level. They are concerned about measurement error in prices. Our main result is robust to price measurement error. We emphasize that adding an additional source of measurement error would make the rationalizability notion less demanding. For our main finding, if we rationalize the single-individual data set without general cases of measurement error, then it will be rationalized if we allow for (non-systematic) measurement error in prices.

7. Recoverability and Counterfactuals

Our general methodology allows us to answer important questions about the recoverability of, and counterfactual predictions for, different objects of a model of interest. In section 7.1 we start by showing how to recover different quantities of interest (e.g., average true consumption at a given $t = \tau$, the support of the discount factor, or the expected value of the random discount factor) from the s/m-rationalizable data set. In section 7.2 we then demonstrate how to make out-of-sample predictions for expected consumption in a way that is analogous to Blundell et al. (2014). In the presence of measurement error, distributional information about the primitives of the model of interest is inevitably lost. Hence, we cannot apply the traditional approach proposed by Varian (1982) to recover preferences and to do counterfactual analysis on an individual basis. Instead, we use this section to pose questions about the primitives of the model at the level of the population.

7.1. Recoverability

Recall that $\mathbf{x}$ and $\mathbf{e}$ denote respectively the observed quantities and the latent random objects. Suppose that there is a finite-dimensional parameter of interest $\theta_0 \in \Theta$, where Θ is a compact subset of the Euclidean space. The parameter of interest is related to the model via the known function $g_R : X \times E \times \Theta \to \mathbb{R}^{d_R}$ such that

$$\mathbb{E}_{\mu \times \pi_0} \left[g_R(\mathbf{x}, \mathbf{e}; \theta_0) \right] = 0.$$

Given function g_R, the revealed-preference inequality restrictions g_I, and the mean budget

as moments or quantiles. We discuss this more extensively in the recoverability section (section 7).

constraints g_B, we can define a new set of MEM as follows:

$$\tilde{h}_{BR}(x;\theta,\gamma)=\frac{\int_{e\in E|X} g_{BR}(x,e;\theta)\exp(\gamma' g_{BR}(x,e;\theta))\mathbb{1}\left(\, g_I(x,e)=0\,\right)d\eta(e|x)}{\int_{e\in E|X}\exp(\gamma' g_{BR}(x,e;\theta))\mathbb{1}\left(\, g_I(x,e)=0\,\right)d\eta(e|x)},$$

where $g_{BR}(x,e;\theta)=(g_B(x,e)',g_R(x,e;\theta)')'$.

The function g_R can take different forms depending on the different questions the user wants to answer. We provide some examples here.

Example 4. (Expected True Consumption/Expected True Consumption Change) If θ_0 is the expected true consumption at $t=\tau$, then $g_R(x,e;\theta_0)=c_\tau-w_\tau-\theta_0$. If θ_0 is an expected difference in true consumption at $t=\tau+1$ and $t=\tau$, then $g_R(x,e;\theta_0)=c_{\tau+1}-w_{\tau+1}-c_\tau+w_\tau-\theta_0$.

The user may also be interested in testing the joint null hypothesis that (i) the consumer is s/ED-rationalizable and (ii) the random discount factor distribution has certain properties.

Example 5. (Average Random Discount Factor) The user may be interested in testing whether the average value of the random discount factor is equal to a certain fixed value, in which case $g_R(x,e;\theta_0)=\mathbf{d}-\theta_0$.

In addition, our framework allows us to have, as a special case, latent random variables with flexible support (that we consider in our empirical application).

Example 6. (Support of the Random Discount Factor) The user may be interested in whether the random time-discount factor $\mathbf{d}=\boldsymbol{\lambda}_t^{-1/t}$ has a support on or inside $[\theta_{01},\theta_{02}]\subseteq(0,1]$. Then, for $\theta_0=(\theta_{01},\theta_{02})'$, one can define $g_R(x,e;\theta_0)=\mathbb{1}\left(\theta_{01}\leq\lambda_1^{-1}\leq\theta_{02}\right)-1$.

As in section 5.4 we can define the sample analogues of the new MEM, the MEM variance matrix, and the test statistic, as follows:

$$\begin{aligned}\hat{\tilde{h}}_{BR}(\theta,\gamma)&=\frac{1}{n}\sum_{i=1}^{n}\tilde{h}_{BR}(\mathbf{x}_i;\theta,\gamma);\\ \hat{\tilde{\Omega}}_{BR}(\theta,\gamma)&=\frac{1}{n}\sum_{i=1}^{n}\tilde{h}_{BR}(\mathbf{x}_i;\theta,\gamma)\tilde{h}_{BR}(\mathbf{x}_i;\theta,\gamma)'-\hat{\tilde{h}}_{BR}(\theta,\gamma)\hat{\tilde{h}}_{BR}(\theta,\gamma)';\\ \mathrm{TS}_n(\theta)&=n\inf_{\gamma\in\mathbb{R}^{q+d_R}}\hat{\tilde{h}}_{BR}(\theta,\gamma)'\hat{\tilde{\Omega}}^{-}_{BR}(\theta,\gamma)\hat{\tilde{h}}_{BR}(\theta,\gamma).\end{aligned}$$

Under assumptions similar to those for Theorem 5, the confidence set for θ_0 can be obtained by inverting $\mathrm{TS}_n(\theta_0)$. That is, the $(1-\alpha)$-confidence set for θ_0 is

$$\{\theta_0\in\Theta:\mathrm{TS}_n(\theta_0)\leq\chi^2_{q+d_R,1-\alpha}\},$$

where $\chi^2_{q+d_R,1-\alpha}$ denotes the $(1-\alpha)$ quantile of the χ^2 distribution with $(q+d_R)$ degrees of freedom $(\chi^2_{q+d_R})$.[33] Note that we do not pretest for s/m-rationalizability in order to construct the

[33]If g_R imposes restrictions only on the support of the latent $\mathbf{e}$, then one can use less conservative critical values of χ^2_q.

confidence set for θ_0. If the data set is not s/m-rationalizable, then the confidence set will be empty asymptotically.

7.2. Counterfactual Out-of-Sample Predictions

We consider a counterfactual situation in which the user is given an out-of-sample deterministic effective price vector $\rho_{T+1} \in \mathbb{R}^L_{++}$, and she then asks two related questions. First, the user wants to know if there exists a counterfactual random consumption vector $\mathbf{c}_{T+1}$ with support $C_{T+1} \subseteq \mathbb{R}^L_+ \setminus \{0\}$, such that the augmented random array $\{(\boldsymbol{\rho}_t, \mathbf{c}_t)_{t\in\mathcal{T}}, (\rho_{T+1}, \mathbf{c}_{T+1})\}$ is approximately s/m-rationalizable (such that $\mathbf{w}_{T+1} = 0$a.s.). (The user is given the support of the extended marginal utility of income $(\boldsymbol{\lambda}_t)_{t\in\mathcal{T}\cup\{T+1\}}$, $\Lambda_{T+1} \subseteq \mathbb{R}^{|\mathcal{T}|+1}_{++}$.)

Second, if the answer to the first question is affirmative, then the user will be interested in constructing confidence sets for some counterfactual finite-dimensional parameter $\theta_0 \in \Theta$. The parameter θ_0 satisfies the user-specified moment condition

$$\mathbb{E}\left[g_C(\mathbf{x}, \mathbf{c_{T+1}}; \rho_{T+1}, \theta_0)\right] = 0.$$

Both questions can be answered simultaneously with our testing methodology in the presence of measurement error. Observe that the answer to the first question is negative if the random array $\mathbf{x}$ is not s/m-rationalizable. In contrast, if the random array $\mathbf{x}$ is s/m-rationalizable and the first question is answered affirmatively, the second question becomes relevant. In the latter case, the counterfactual exercise is equivalent to testing the joint null hypothesis that the counterfactual price/consumption distribution is simultaneously compatible with s/m-rationalizability and the user-specified moment condition. Formally, to answer both questions, we define what it means for a random array $\mathbf{x}$ to be counterfactually rationalizable in the presence of measurement error (C/m-rationalizability) for a given ρ_{T+1}, θ_0, and g_C.

Definition 8. (C/m-rationalizability) For a given ρ_{T+1} and θ_0, a random array $\mathbf{x}$ is said to be C/m-rationalizable if there exist a tuple

$$(\mathbf{u}, (\boldsymbol{\lambda}_t)_{t\in\mathcal{T}\cup\{T+1\}}, (\boldsymbol{w}_t)_{t\in\mathcal{T}\cup\{T+1\}}, \mathbf{c}_{T+1}),$$

such that:

(i) $\mathbf{u}$ is a random, concave, locally nonsatiated, and continuous utility function;

(ii) $(\boldsymbol{\lambda}_t)_{t\in\mathcal{T}\cup\{T+1\}}$ is a positive random vector, interpreted as the marginal utility of income, supported on or inside a given $\Lambda_{T+1} \subseteq \mathbb{R}^{|\mathcal{T}|+1}_{++}$;

(iii) $(\boldsymbol{w}_t)_{t\in\mathcal{T}\cup\{T+1\}}$ is a random array that defines $\mathbf{c}^*_t = \mathbf{c}_t - \mathbf{w}_t$, where $\mathbf{c}^*_t$ is a random (true) consumption vector supported on or inside $C^*_t \subseteq \mathbb{R}^L_+ \setminus \{0\}$ for $t \in \mathcal{T}$, and $\mathbf{w}_{T+1} = 0$ a.s.;

(iv) $\mathbf{c}_{T+1}$ is a random (counterfactual) consumption vector supported on or inside $C_{T+1} \subseteq \mathbb{R}^L_+ \setminus \{0\}$;

(v) $\nabla \mathbf{u}(\mathbf{c}_t^*) \leq \boldsymbol{\lambda}_t \boldsymbol{\rho}_t$ a.s. for $t \in \mathcal{T}$, and $\nabla \mathbf{u}(\mathbf{c}_{T+1}^*) \leq \boldsymbol{\lambda}_{T+1} \rho_{T+1}$ a.s.;

(vi) for every $j = 1, \ldots, L$ and $t \in \mathcal{T} \cup \{T+1\}$, it must be that $\mathbb{P}\left(\mathbf{c}_{t,j}^* \neq 0, \nabla \boldsymbol{u}(\mathbf{c}_t^*)_j < \boldsymbol{\lambda}_t \boldsymbol{\rho}_{t,j}\right) = 0$, where $c_{t,j}^*$, $\rho_{t,j}$, and $\nabla u(c_t^*)_j$ denote the j-th components of c_t^*, ρ_t and $\nabla u(c_t^*)$, respectively;

(vii) $\mathbb{E}\left[g_C(\mathbf{x}, \mathbf{c}_{T+1}; \rho_{T+1}, \theta_0)\right] = 0$ with $\theta_0 \in \Theta$.

Observe that if a random array $\mathbf{x}$ is C/m-rationalizable for a given ρ_{T+1} and θ_0, then it is also s/m-rationalizable. However, the s/m-rationalizability of $\mathbf{x}$ does not imply that $\mathbf{x}$ is C/m-rationalizable.

It is straightforward to formulate additional moment conditions that correspond to C/m-rationalizability. To do so, we define the augmented latent random array $\tilde{\mathbf{e}} = (\mathbf{e}', \mathbf{v}_{T+1}, \boldsymbol{\lambda}_{T+1}, \mathbf{c}_{T+1}')'$ with known conditional support on or inside $\tilde{E}|X$. Note that the random array $\mathbf{e}$, defined earlier, is supported on or inside $E|X$, because $(\boldsymbol{\lambda}_t)_{t\in\mathcal{T}\cup\{T+1\}}$, is supported or on inside the known support $\Lambda_{T+1} \subseteq \mathbb{R}_{++}^{|\mathcal{T}|+1}$ (which is compatible with Λ). Also, $\mathbf{v}_{T+1}$ is a positive random variable, and $\mathbf{c}_{T+1}$ is supported on or inside C_{T+1}. We then can define additional moments as follows:

$$
\begin{aligned}
g_{O,t,1}(\mathbf{x}, \tilde{\mathbf{e}}; \rho_{T+1}) &= \mathbb{1}\left(\mathbf{v}_{T+1} - \mathbf{v}_t - \lambda_{T+1}\boldsymbol{\rho}_{T+1}'[\mathbf{c}_{T+1} - \mathbf{c}_t + \mathbf{w}_t] \geq 0\right) - 1, \quad t \in \mathcal{T};\\
g_{O,t,2}(\mathbf{x}, \tilde{\mathbf{e}}; \rho_{T+1}) &= \mathbb{1}\left(\mathbf{v}_t - \mathbf{v}_{T+1} - \boldsymbol{\lambda}_t\boldsymbol{\rho}_t'[\mathbf{c}_t - \mathbf{w}_t - \mathbf{c}_{T+1}] \geq 0\right) - 1, \quad t \in \mathcal{T};\\
g_O(\mathbf{x}, \tilde{\mathbf{e}}; \rho_{T+1}) &= (g_{O,t,j}(\tilde{\mathbf{x}}, \tilde{\mathbf{e}}))_{t\in\mathcal{T};j\in\{1,2\}}.
\end{aligned}
$$

Given the original inequality conditions (g_I), the mean budget-neutrality conditions (g_B), and the new counterfactual conditions (g_C), we can now define the new system of moments (g):

$$
\begin{aligned}
g_{I,O}(\mathbf{x}, \tilde{\mathbf{e}}; \rho_{T+1}) &= (g_I(\mathbf{x}, \mathbf{e})', g_O(\mathbf{x}, \tilde{\mathbf{e}}, \rho_{T+1})')';\\
g_{C,B}(\mathbf{x}, \tilde{\mathbf{e}}; \rho_{T+1}, \theta) &= (g_C(\mathbf{x}, \mathbf{c}_{T+1}, \rho_{T+1}, \theta)', g_B(\mathbf{x}, \mathbf{e})')';\\
g(\mathbf{x}, \tilde{\mathbf{e}}; \rho_{T+1}, \theta) &= (g_{I,O}(\mathbf{x}, \tilde{\mathbf{e}}, \rho_{T+1})', g_{C,B}(\mathbf{x}, \tilde{\mathbf{e}}, \rho_{T+1}, \theta)')'.
\end{aligned}
$$

Note that because of the addition of extra latent variables the number of the inequality conditions has increased by $2\mathcal{T}$, while the number of the mean budget-neutrality conditions has not changed. We let $g_C(x, c_{T+1}, \rho_{T+1}, \theta) \in \mathbb{R}^{q_c}$. We can now define a new set of MEM as follows:

$$
\tilde{h}_{C,B}(x; \rho_{T+1}, \theta, \gamma) = \frac{\int_{\tilde{e}\in\tilde{E}|X} g_{C,B}(x, \tilde{e}, \rho_{T+1}, \theta) \exp(\gamma' g_{C,B}(x, \tilde{e}, \rho_{T+1}, \theta)) \mathbb{1}\left(g_{I,O}(x, \tilde{e}, \rho_{T+1}) = 0\right) d\eta(\tilde{e}|x)}{\int_{\tilde{e}\in\tilde{E}|X} \exp(\gamma' g_{C,B}(x, \tilde{e}, \rho_{T+1}, \theta)) \mathbb{1}\left(g_{I,O}(x, \tilde{e}, \rho_{T+1}) = 0\right) d\eta(\tilde{e}|x)},
$$

with $\gamma \in \mathbb{R}^{q+q_c}$.

Now we are ready to present the main result of this section that is an extension of Theorem 3.

Theorem 6. *The following are equivalent:*

(i) For a given ρ_{T+1} and θ_0, a random array $\mathbf{x} = (\boldsymbol{\rho}_t, \mathbf{c}_t)_{t\in\mathcal{T}}$ is approximately C/m-*rationalizable such that Assumptions 1-3 hold.*

(ii)

$$\inf_{\gamma \in \mathbb{R}^{q+q_c}} \left\| \mathbb{E}_{\pi_0} \left[\tilde{h}_{C,B}(\mathbf{x}; \rho_{T+1}, \theta_0, \gamma) \right] \right\| = 0,$$

where $\pi_0 \in \boldsymbol{P}_{\mathcal{X}}$ is the observed distribution of $\mathbf{x}$.

If, for a given ρ_{T+1} and θ_0, the random array $\mathbf{x}$ cannot be C/m-rationalizable, then the second question we posed at the beginning of the section must be answered in the negative. Therefore, this formulation allows us to build confidence sets for θ_0 given ρ_{T+1}, using the results presented in section 5.4 (after imposing regularity conditions on the counterfactual moments $g_C(\cdot)$).

Example 7. (Average Varian Support Set) We consider a moment

$$g_C(\mathbf{x}, \mathbf{c}_{T+1}; \rho_{T+1}, \theta_0) = \mathbf{c}_{T+1} - \theta_0,$$

with $\theta_0 \in \mathbb{R}^L_+ \setminus \{0\}$ as a hypothesized average-demand vector. Given this, we define the Average Varian Support Set as

$$\{\theta \in \mathbb{R}^L_+ \setminus \{0\} | \inf_{\gamma \in \mathbb{R}^{q+q_c}} \left\| \mathbb{E}_{\pi_0} \left[\tilde{h}_{C,B}(\mathbf{x}; \rho_{T+1}, \theta, \gamma) \right] \right\| = 0\}.$$

This set describes the bounds of the average demand given ρ_{T+1}, that is compatible with the s/m-rationalizability of the random array $\mathbf{x}$.

Example 8. (Quantile Varian Support Set) For the case of s/R-rationalizability, we can consider the following moment condition:

$$g_C(\mathbf{x}, \mathbf{c}_{T+1}; \rho_{T+1}, \theta) = \mathbb{1}\left(\rho'_{T+1}\mathbf{c}_{T+1} \leq \bar{e}_c \right) - \alpha,$$

where $\theta = (\bar{e}_c, \alpha)' \in \mathbb{R}_{++} \times [0, 1]$, $\bar{e}_c$ is a fixed α-quantile of the counterfactual expenditure distribution. Next we can define the α-quantile Varian Support Set:

$$\left\{ c \in \mathbb{R}^L_+ \setminus \{0\} \ : \ \rho'_{T+1} c = \bar{e}_c, \inf_{\gamma \in \mathbb{R}^{q+q_c}} \| \mathbb{E}_{\pi_0} \left[h_{C,B}(\mathbf{x}; \rho_{T+1}, \theta, \gamma) \right] \| = 0 \right\}.$$

This set describes the bounds of the counterfactual demand for a given ρ_{T+1} and α-quantile of $\mathbf{u}(\mathbf{c}_{T+1})$ that is compatible with s/R-rationalizability.

We believe that the Average Varian Support Set and the α-Quantile Varian Support Set are potentially interesting objects for practitioners of the revealed-preference methodology in survey data. In particular, we note that these objects are similar to the demand bounds proposed by Blundell et al. (2003), in the presence of measurement error, but without infinite variation in income. We must emphasize once again that in contrast to Blundell et al. (2014) we do not impose rationalizability but instead test for it.

8. Conclusion

We propose a new stochastic and nonparametric revealed-preference approach (suitable for an environment with measurement error in consumption) that is useful to test for several consumer models that can be characterized by their first-order conditions. In particular, our work can be used (but is not limited) to test for static utility maximization (Afriat (1967)), for quasilinear utility maximization (Brown and Calsamiglia (2007)), and for dynamic rationalizability with exponential discounting (Browning (1989)). We provide Monte Carlo evidence suggesting that the deterministic revealed-preference approach may fail if the observed data has measurement error in consumption. The methodology presented here is able to outperform the deterministic approach in this particular environment.

We apply our methodology to a widely used consumption panel household survey, and we find robust evidence against dynamic rationalizability with exponential discounting in the case of couples' households. More surprisingly, we cannot reject the hypothesis that the data set is statistically rationalizable by an exponential discounting consumer model for the case of single-individual households. This finding goes against the conclusion reached after applying the deterministic revealed-preference approach (Browning (1989)) to the same data set. We believe this is convincing evidence that taking into account measurement error is crucial for using the revealed-preference methodology in survey data.

Acceptance of the exponential discounting model means that in some situations, such as in our application, consumers may behave in a time-consistent manner. This may happen when the decision is made frequently, which can lead to developing expertise. More important, the decisions that we study, in contrast to the low-stakes decisions in laboratory situations, involve moderate stakes. The consumers may be more likely to exert self-control and behave in a time-consistent manner when making decisions about purchasing food and services. In our application the consumers allocate their budget shares among different good-categories, such as food, clothing, and transportation. Arguably, this type of decision is both repetitive and important enough for the consumers to be time-consistent. In contrast, for the case of couples' households we reject the unitary exponential discounting model, while we fail to reject the collective exponential discounting model with full efficiency. Our results when compared with the single household evidence suggest that time-inconsistencies in the consumption behavior in the couples' case arise due to preference aggregation. We leave as an open question the empirical and theoretical comparison of time-inconsistent models such as hyperbolic discounting versus models of collective decision-making with preference aggregation, in the presence of measurement error.

References

Adams, A., Cherchye, L., De Rock, B., and Verriest, E. (2014). Consume Now or Later? Time Inconsistency, Collective Choice, and Revealed Preference. *American Economic Review*, 104(12):4147–4183.

Afriat, S. N. (1967). The construction of utility functions from expenditure data. *International economic review*, 8(1):67–77.

Ai, C. and Chen, X. (2003). Efficient estimation of models with conditional moment restrictions containing unknown functions. *Econometrica*, 71(6):1795–1843.

Andreoni, J. and Sprenger, C. (2012). Estimating time preferences from convex budgets. *The American Economic Review*, 102(7):3333–3356.

Beatty, T. K. and Crawford, I. A. (2011). How demanding is the revealed preference approach to demand? *The American Economic Review*, 101(6):2782–2795.

Blow, L., Browning, M., and Crawford, I. (2013). Never Mind the Hyperbolics: Nonparametric Analysis of Time-Inconsistent Preferences. *Unpublished manuscript.*

Blow, L., Browning, M., and Crawford, I. (2017). Nonparametric analysis of time-inconsistent preferences. *Mimeo.*

Blundell, R., Kristensen, D., and Matzkin, R. (2014). Bounding quantile demand functions using revealed preference inequalities. *Journal of Econometrics*, 179(2):112–127.

Blundell, R. W., Browning, M., and Crawford, I. A. (2003). Nonparametric engel curves and revealed preference. *Econometrica*, 71(1):205–240.

Boccardi, M. J. (2016). Predictive ability and the fit-power trade-off in theories of consumer behavior. *Mimeo.*

Brown, D. J. and Calsamiglia, C. (2007). The nonparametric approach to applied welfare analysis. *Economic Theory*, 31(1):183–188.

Browning, M. (1989). A nonparametric test of the life-cycle rational expections hypothesis. *International Economic Review*, pages 979–992.

Campbell, J. Y. and Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of political Economy*, 107(2):205–251.

Carroll, C. D., Crossley, T. F., and Sabelhaus, J. (2014). Introduction to" improving the measurement of consumer expenditures". In *Improving the Measurement of Consumer Expenditures*, pages 1–20. University of Chicago Press.

Chapman, D. A. (1997). Approximating the asset pricing kernel. *The Journal of Finance*, 52(4):1383–1410.

Chen, X., Chernozhukov, V., Lee, S., and Newey, W. K. (2014). Local identification of nonparametric and semiparametric models. *Econometrica*, 82(2):785–809.

Chen, X. and Ludvigson, S. C. (2009). Land of addicts? an empirical investigation of habit-based asset pricing models. *Journal of Applied Econometrics*, 24(7):1057–1093.

Cherchye, L., Demuynck, T., De Rock, B., Vermeulen, F., et al. (2017). Household consumption when the marriage is stable. *American Economic Review*, 107(6):1507–1534.

Choi, S., Kariv, S., Müller, W., and Silverman, D. (2014). Who is (more) rational? *The American Economic Review*, 104(6):1518–1550.

Crawford, I. (2010). Habits revealed. *The Review of Economic Studies*, 77(4):1382–1402.

Darolles, S., Fan, Y., Florens, J.-P., and Renault, E. (2011). Nonparametric instrumental regression. *Econometrica*, 79(5):1541–1565.

Deb, R., Kitamura, Y., Quah, J. K.-H., and Stoye, J. (2017). Revealed price preference: Theory and stochastic testing.

DellaVigna, S. and Malmendier, U. (2006). Paying not to go to the gym. *The American Economic Review*, pages 694–719.

Dette, H., Hoderlein, S., and Neumeyer, N. (2016). Testing multivariate economic restrictions using quantiles: the example of slutsky negative semidefiniteness. *Journal of Econometrics*, 191(1):129–144.

Dunn, K. B. and Singleton, K. J. (1986). Modeling the term structure of interest rates under non-separable utility and durability of goods. *Journal of Financial Economics*, 17(1):27–55.

Echenique, F., Imai, T., and Saito, K. (2014). Testable Implications of Quasi-Hyperbolic and Exponential Time Discounting.

Ekeland, I., Galichon, A., and Henry, M. (2010). Optimal transportation and the falsifiability of incompletely specified economic models. *Economic Theory*, 42(2):355–374.

Emiris, I. Z. and Fisikopoulos, V. (2013). Algorithms for volume approximation of convex bodies. Technical report, Technical Report CGL-TR-76, NKUA.

Escanciano, J. C., Hoderlein, S., Lewbel, A., Linton, O., and Srisuma, S. (2016). Nonparametric euler equation identification and estimation. *Working Paper*.

Forges, F. and Minelli, E. (2009). Afriat's theorem for general budget sets. *Journal of Economic Theory*, 144(1):135–145.

Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature*, 40(2):351–401.

Galichon, A. and Henry, M. (2013). Dilation bootstrap. *Journal of Econometrics*, 177(1):109 – 115.

Gallant, A. R. and Tauchen, G. (1989). Seminonparametric estimation of conditionally constrained heterogeneous processes: Asset pricing applications. *Econometrica: Journal of the Econometric Society*, pages 1091–1120.

Gillen, B., Snowberg, E., and Yariv, L. (2017). Experimenting with measurement error: techniques with applications to the caltech cohort study. Technical report, National Bureau of Economic Research.

Gross, J. (1995). Testing data for consistency with revealed preference. *The Review of Economics and Statistics*, pages 701–710.

Halevy, Y. (2015). Time consistency: Stationarity and time invariance. *Econometrica*, 83(1):335–352.

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence. *Journal of political economy*, 86(6):971–987.

Hansen, L. P. and Singleton, K. J. (1982). Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica: Journal of the Econometric Society*, pages 1269–1286.

Hjertstrand, P. (2013). A simple method to account for measurement errors in revealed preference tests.

Hoderlein, S., Vanhems, A., et al. (2013). Estimating the distribution of welfare effects using quantiles. *Work. Pap., Dep. Econ., Boston Coll.*

Jackson, M. O. and Yariv, L. (2015). Collective dynamic choice: the necessity of time inconsistency. *American Economic Journal: Microeconomics*, 7(4):150–178.

Jerison, D. and Jerison, M. (1994). Commodity aggregation and slutsky asymmetry. In *Models and Measurement of Welfare and Inequality*, pages 752–764. Springer.

Kitamura, Y. and Stoye, J. (2016). Nonparametric analysis of random utility models. Technical report, cemmap working paper, Centre for Microdata Methods and Practice.

Lewbel, A. and Pendakur, K. (2017). Unobserved preference heterogeneity in demand using generalized random coefficients. *Journal of Political Economy*, 125(4):1100–1148.

Marschak, J. (1974). *Economic Measurements for Policy and Prediction*, pages 293–322. Springer Netherlands, Dordrecht.

Mastrobuoni, G. and Rivers, D. A. (2016). Criminal discount factors and deterrence.

Mathiowetz, N., Brown, C., and Bound, J. (2002). Measurement error in surveys of the low-income population. *Studies of welfare populations: Data collection and research issues*, pages 157–194.

Montiel Olea, J. L. and Strzalecki, T. (2014). Axiomatization and Measurement of Quasi-hyperbolic Discounting. *Quarterly Journal of Economics*, 129:1449–1499.

Rockafellar, R. T. (1970). *Convex analysis.* Princeton university press.

Schennach, S. M. (2014). Entropic latent variable integration via simulation. *Econometrica*, 82(1):345–385.

Sulka, T. (2017). Cognitive costs and corner solutions in retirement planning. *Mimeo.*

Tsur, Y. (1989). On testing for revealed preference conditions. *Economics Letters*, 31(4):359–362.

Varian, H. R. (1982). The nonparametric approach to demand analysis. *Econometrica: Journal of the Econometric Society*, pages 945–973.

Varian, H. R. (1984). The nonparametric approach to production analysis. *Econometrica: Journal of the Econometric Society*, pages 579–597.

Varian, H. R. (1985). Non-parametric analysis of optimizing behavior with measurement error. *Journal of Econometrics*, 30(1-2):445–458.

Varian, H. R. (1990). Goodness-of-fit in optimizing models. *Journal of Econometrics*, 46(1-2):125–140.

Weyl, E. G. (2009). Slutsky meets marschak: The first-order identification of multi-product production.

9. Appendix

9.1. Proof of Lemma 2

Proof. First we establish that (i) implies (ii). If the random array $(\boldsymbol{\rho}_t, \mathbf{c}_t^*)_{t\in\mathcal{T}}$ is s/m-rationalizable, by concavity of $\mathbf{u}(\cdot)$, with probability 1 for any s, t and some $\xi \in \nabla \mathbf{u}(\mathbf{c}_t)$

$$\mathbf{u}(\mathbf{c}_s^*) - \mathbf{u}(\mathbf{c}_t^*) \le \xi'(\mathbf{c}_s^* - \mathbf{c}_t^*),$$
$$\xi \le \boldsymbol{\lambda}_t \boldsymbol{\rho}_t.$$

Let $\mathbf{N}$ be a random set of indices such that $\boldsymbol{\lambda}_t\boldsymbol{\rho}_{ti} = \xi_i$ for every $i \in \mathbf{N}$. Hence, $\boldsymbol{\lambda}_t\boldsymbol{\rho}_{ti} \geq \xi_i$ for every $i \notin \mathbf{N}$ with probability 1. As a result, $\mathbf{c}^*_{ti}$ has to be a corner solution for every $i \notin \mathbf{N}$. That is, $\mathbf{c}^*_{ti} = 0$. Thus, with probability 1,

$$\mathbf{u}(\mathbf{c}^*_s) - \mathbf{u}(\mathbf{c}^*_t) \leq \xi'(\mathbf{c}^*_s - \mathbf{c}^*_t) = \sum_{i\in\mathbf{N}} \xi_i(\mathbf{c}^*_{si} - \mathbf{c}^*_{ti}) + \sum_{i\notin\mathbf{N}} \xi_i \mathbf{c}^*_{si} =$$
$$= \sum_{i\in\mathbf{N}} \boldsymbol{\lambda}_t\boldsymbol{\rho}_{ti}(\mathbf{c}^*_{si} - \mathbf{c}^*_{ti}) + \sum_{i\notin\mathbf{N}} \xi_i\mathbf{c}^*_{si} \leq \sum_{i\in\mathbf{N}} \boldsymbol{\lambda}_t\boldsymbol{\rho}_{ti}(\mathbf{c}^*_{si} - \mathbf{c}^*_{ti}) + \sum_{i\notin\mathbf{N}} \boldsymbol{\lambda}_t\boldsymbol{\rho}_{ti}\mathbf{c}^*_{si},$$

where the last inequality follows from $\mathbf{c}_s$ being nonnegative. As a result, with probability 1,

$$\forall s, t \in \mathcal{T} \;:\; \mathbf{u}(\mathbf{c}^*_t) - \mathbf{u}(\mathbf{c}^*_s) \geq \boldsymbol{\lambda}_t\boldsymbol{\rho}'_t[\mathbf{c}^*_t - \mathbf{c}^*_s].$$

For any $\epsilon > 0$, we let $\mathbf{v}_t = \mathbf{u}(\mathbf{c}^*_t) - \min_{s\in\mathcal{T}} \mathbf{u}(\mathbf{c}^*_s) + \epsilon$ a.s., for all $t \in \mathcal{T}$. The well-defined positive random vector $(\mathbf{v}_t)_{t\in\mathcal{T}}$ together with $(\boldsymbol{\lambda}_t)_{t\in\mathcal{T}}$ satisfies the inequalities in (ii).

Now, we want to prove that (ii) implies (i). The result follows from Theorem 24.8 in Rockafellar (1970). For completeness of the proof we repeat the arguments of Theorem 24.8 in Rockafellar (1970). For a finite $m \in \mathbb{N}$, let $\mathbf{t} = \{t_i\}_{i=1}^m$, $t_i \in \mathcal{T}$, be a finite set of indices such that for a fixed $\hat{t} \in \mathcal{T}$, $c^*_{t_1} = c^*_{\hat{t}}$. Let $\mathcal{I}$ be the collection of all such indices (i.e., $\mathbf{t} \in \mathcal{I}$). Define

$$\mathbf{u}(c^*) = \inf_{\mathbf{t}\in\mathcal{I}}\{\boldsymbol{\lambda}_{t_1}\boldsymbol{\rho}'_{t_1}(\mathbf{c}^*_{t_2} - \mathbf{c}^*_{t_1}) + \cdots + \boldsymbol{\lambda}_{t_m}\boldsymbol{\rho}'_{t_m}(c^* - \mathbf{c}^*_{t_m})\}.$$

With probability 1, the random function $\mathbf{u}(\cdot)$ is well-defined, concave, locally nonsatiated, and continuous, since it is a pointwise minimum of a finite set of affine functions for every m. Moreover, the infimum in $\mathcal{I}$ is attained and achieved at a set of indices without repetitions (this is a consequence of (ii)). Indeed, under (ii), for any finite m, $\{t_i\}_{i=1}^m$ and $\mathbf{c}^*_s$, $s \in \mathcal{T}$, with probability 1,

$$\boldsymbol{\lambda}_{t_1}\boldsymbol{\rho}'_{t_1}(\mathbf{c}^*_{t_2} - \mathbf{c}^*_{t_1}) + \cdots + \boldsymbol{\lambda}_{t_m}\boldsymbol{\rho}'_{t_m}(\mathbf{c}^*_s - \mathbf{c}^*_{t_m}) + \boldsymbol{\lambda}_s\boldsymbol{\rho}'_s(\mathbf{c}^*_{t_1} - \mathbf{c}^*_s) \geq$$
$$\mathbf{v}_{t_2} - \mathbf{v}_{t_1} + \mathbf{v}_{t_3} - \mathbf{v}_{t_2} + \cdots + \mathbf{v}_s - \mathbf{v}_{t_m} + \mathbf{v}_{t_1} - \mathbf{v}_s = 0.$$

Thus,

$$\mathbf{u}(\mathbf{c}^*_s) \geq \boldsymbol{\lambda}_s\boldsymbol{\rho}'_s(\mathbf{c}^*_s - \mathbf{c}^*_{t_1}) > -\infty$$

with probability 1. (In particular, $u(c^*_{\hat{t}}) = 0$.)

It is left to show that, with probability 1, $\boldsymbol{\lambda}_t\boldsymbol{\rho}_t \in \nabla\mathbf{u}(\mathbf{c}^*_t)$ for all $t \in \mathcal{T}$. Fix some $t \in \mathcal{T}$ and $\delta > 0$. By the definition of $\mathbf{u}(\cdot)$, there exists some $\{t_i\}_{i=1}^m$ such that, with probability 1, $\mathbf{u}(\mathbf{c}^*_t) + \delta > \boldsymbol{\lambda}_{t_1}\boldsymbol{\rho}'_{t_1}(\mathbf{c}^*_{t_2} - \mathbf{c}^*_{t_1}) + \cdots + \boldsymbol{\lambda}_{t_m}\boldsymbol{\rho}'_{t_m}(\mathbf{c}^*_t - \mathbf{c}^*_{t_m}) \geq \mathbf{u}(\mathbf{c}^*_t)$. Again, by the definition of $\mathbf{u}(\cdot)$, for any c^*

$$\boldsymbol{\lambda}_{t_1}\boldsymbol{\rho}'_{t_1}(\mathbf{c}^*_{t_2} - \mathbf{c}^*_{t_1}) + \cdots + \boldsymbol{\lambda}_{t_m}\boldsymbol{\rho}'_{t_m}(\mathbf{c}^*_t - \mathbf{c}^*_{t_m}) + \boldsymbol{\lambda}_t\boldsymbol{\rho}'_t(c^* - \mathbf{c}^*_t) \geq \mathbf{u}(c^*).$$

Hence,

$$\mathbf{u}(\mathbf{c}^*_t) + \delta + \boldsymbol{\lambda}_t\boldsymbol{\rho}'_t(c^* - \mathbf{c}^*_t) > \mathbf{u}(c^*).$$

Since the choice of δ, t and c^* was arbitrary, $\boldsymbol{\lambda}_t\boldsymbol{\rho}_t \in \nabla\mathbf{u}(\mathbf{c}^*_t)$ for all $t \in \mathcal{T}$.

$\blacksquare$

9.2. Proof of Theorem 3

Proof. The result is a direct application of Theorem 2, and Theorem 2.1 in Schennach (2014). For completeness of the proof we present Theorem 2.1 in Schennach (2014) using our notation below.

Theorem. *(Theorem 2.1, Schennach (2014)) Assume that the marginal distribution of x is supported on some set $X \subseteq \mathbb{R}^{d_x}$, while the distribution of $\mathbf{e}$ conditional on $\mathbf{x} = x$ is supported on or inside the set $E \subseteq \mathbb{R}^{d_e}$ for any $x \in X$. Let h, g and η satisfy Definition 6. Then*

$$\inf_{\mu \in \mathcal{P}_{E|X}} \|\mathbb{E}_{\mu \times \pi_0}[\, g(\mathbf{x}, \mathbf{e})\,]\| = 0 \iff \inf_{\gamma \in \mathbb{R}^{k+q}} \|\mathbb{E}_{\pi_0}[\, h(\mathbf{x}; \gamma)\,]\| = 0,$$

where $\pi_0 \in \boldsymbol{P}_{\mathcal{X}}$ is the observed distribution of $\mathbf{x}$.

$\blacksquare$

9.3. Proof of Theorem 4

Recall that the first $k = |\mathcal{T}|^2 - |\mathcal{T}|$ moments correspond to the inequality conditions, and the last $q = |\mathcal{T}|$ moments correspond to the budget neutrality conditions. Let $\gamma_I = (\gamma_j)_{j=1,\dots,k}$, $g_I = (g_j)_{j=1,\dots,k}$, $\gamma_B = (\gamma_j)_{j=k+1,\dots,k+q}$, and $g_B = (g_j)_{j=k+1,\dots,k+q}$ be sub-vectors of γ and g that correspond to inequality and the mean budget-neutrality conditions respectively.

Proof. Step 1. Take a sequence $\{\gamma_{I,l}\}_{l=1}^{+\infty}$ such that every component of $\gamma_{I,l}$ diverges to $+\infty$. Note that since g_I takes values in $\{-1, 0\}^k$,

$$\sup_{x,e} \left| \exp(\gamma_{I,l}' g_I(x,e)) - \mathbb{1}\,(\, g_I(x,e) = 0\,) \right| \le \exp(-\min_{i=1,\dots,k} \gamma_{I,l,i}) \to_{l \to +\infty} 0,$$

where $\gamma_{I,l,i}$ is the i-th component of $\gamma_{I,l}$. Hence, for any function $f \in L^1(\eta(\cdot|x))$

$$\begin{aligned}&\left\| \int f(e) \exp(\gamma_{I,l}' g_I(x,e)) d\eta(e|x) - \int f(e) \mathbb{1}\,(\, g_I(x,e) = 0\,)\, d\eta(e|x) \right\| \le \\ &\le \exp(-\min_{i=1,\dots,k} \gamma_{I,l,i}) \int \|f(e)\|\, d\eta(e|x) \to_{l \to +\infty} 0.\end{aligned}$$

Hence, the sequence of measures $\exp(\gamma_{I,l}' g_I(x,\cdot)) d\eta(\cdot|x)$ converges to the measure $\mathbb{1}\,(\, g_I(x,\cdot) = 0\,)\, d\eta(\cdot|x)$ in total variation. The later measure is well defined and nontrivial since we assume that $\tilde{E}|X = \{e : \mathbb{1}\,(\, g_I(x,e) = 0\,)\}$ has a positive measure under $\eta(\cdot|x)$. Let $d\tilde{\eta}(\cdot|x)$ denote $\mathbb{1}\,(\, g_I(x,\cdot) = 0\,)\, d\eta(\cdot|x)$.

Step 2. Consider the moment conditions under $d\tilde{\eta}(\cdot|x)$

$$\tilde{h}_B(x;\gamma) = \frac{\int_{e \in E|X} g_B(x,e) \exp(\gamma' g_B(x,e)) d\tilde{\eta}(e|x)}{\int_{e \in E|X} \exp(\gamma' g_B(x,e)) d\tilde{\eta}(e|x)}.$$

Definition 7.(iii) together with Assumption 3 and Step 1 imply that for any compact set $\Gamma \in \mathbb{R}^q$, uniformly in $\gamma_B \in \Gamma$

$$\left\|\mathbb{E}_{\pi_0}\left[h(\mathbf{x};(\gamma'_{I,l},\gamma'_B)')\right]\right\| = \left\|\mathbb{E}_{\pi_0}\left[\tilde{h}_B(\mathbf{x};\gamma_B)\right]\right\| + o(1).$$

Thus, by continuity of h_B in γ_B, when l goes to infinity, we can work with the reduced optimization problem:

$$\inf_{\gamma\in\mathbb{R}^q}\left\|\mathbb{E}_{\pi_0}\left[\tilde{h}_B(\mathbf{x};\gamma)\right]\right\|. \tag{5}$$

Step 3. Note that (5) is equivalent to the optimization problem in Theorem 3. Hence, infimum in (5) is equal to 0 if and only if the data is approximately consistent with model m.

We assumed that every component of g_B takes both positive and negative values on some non-zero measure subsets of $\tilde{E}|X$ (Assumption 2). Hence, following the proof of Theorem 2.1 and Lemma A.1 in Schennach (2014), we can conclude that if infimum in (5) is equal to 0, then it is achieved at some finite and unique $\gamma_{0,B}$. Otherwise, $\|\gamma_B\|$ diverges to infinity. ■

9.4. Proof of Theorem 5

Proof. The result is a direct application of Theorem F.1 in Schennach (2014). For completeness of the proof we present the version of it that is applicable to our setting below.

Theorem. *(Theorem F.1, Schennach (2014)) Let data be i.i.d.. If (i) the set*

$$\Gamma = \{\gamma \in \mathbb{R}^q : \mathbb{E}\left[\left\|\tilde{h}(\mathbf{x},\gamma)\right\|\right] \leq C\}$$

is nonempty for some $C < \infty$; (ii) $\mathbb{E}\left[\left\|\tilde{h}(\mathbf{x},\gamma)\right\|^2\right] < \infty$ for all $\gamma \in \Gamma$, then

$$\lim_{n\to\infty}\mathbb{P}\left(\mathrm{TS}_n > \chi^2_{q,\alpha}\right) \leq \alpha.$$

An i.i.d. sample is assumed. To show the validity of conditions (i) and (ii) note that since $\mathbf{x}$ has a bounded support (by Assumption 3) and $\tilde{\eta}(\cdot|\cdot)$ satisfies conditions of Definition 7.(iii), for any finite γ there exist finite positive constant $C_1(\gamma)$ such that almost surely in $\mathbf{x}$

$$\left\|\tilde{h}(\mathbf{x},\gamma)\right\|^2 \leq C_1(\gamma).$$

Hence, for any nonempty compact set Γ one can take $C = \sup_{\gamma\in\Gamma} C_1(\gamma)$. Together with Assumption 3, the later implies condition (ii). Similarly, one can use C to bound $\mathbb{E}\left[\left\|\tilde{h}(\mathbf{x},\gamma)\right\|\right]$.

Under the alternative hypothesis, $\left\|\hat{\tilde{h}}(\gamma)\right\|$ either converges to a positive constant or diverges to infinity. Thus, since eigenvalues of $\tilde{\Omega}(\gamma)$ are bounded away from zero and are bounded from above the test is consistent. ■

9.5. Proof of Theorem 8

Proof. By Theorem (7) we have that the following inequalities hold a.s.:

$$\mathbf{v}_{t,A} - \mathbf{v}_{s,A} \geq \frac{1}{\mathbf{d}_A^t}[\boldsymbol{\rho}'_{t,I}(\mathbf{c}^*_{t,I} - \mathbf{c}^*_{t,B} - \mathbf{c}^*_{s,I} + \mathbf{c}^*_{s,B}) + \frac{\mathbf{p}_{t,H} - \mathbf{p}_{t,B}}{\prod_{j=1}^t (1+\mathbf{r}_j)}'(\mathbf{c}^*_{t,H} - \mathbf{c}^*_{s,H}) \quad \forall t, s \in \mathcal{T},$$

$$\mathbf{v}_{t,B} - \mathbf{v}_{s,B} \geq \frac{1}{\mathbf{d}_B^t}[\boldsymbol{\rho}'_{t,I}(\mathbf{c}^*_{t,B} - \mathbf{c}^*_{s,B}) + \frac{\mathbf{p}_{t,B}}{\prod_{j=1}^t (1+\mathbf{r}_j)}'(\mathbf{c}^*_{t,H} - \mathbf{c}^*_{s,H}) \quad \forall t, s \in \mathcal{T}.$$

Then we multiply the first inequality by $\mathbf{d}_A^t$, this random variable is positive a.s., so it does not alter the inequalities. We do the same for the second inequality, and multiply it by $\mathbf{d}_B^t$. Then we add-up the two inequalities, to obtain:

$$\mathbf{d}_A^t(\mathbf{v}_{t,A} - \mathbf{v}_{s,A}) + \mathbf{d}_B^t(\mathbf{v}_{t,B} - \mathbf{v}_{s,B}) \geq \boldsymbol{\rho}'_{t,I}(\mathbf{c}^*_t - \mathbf{c}^*_s) \quad \forall t, s \in \mathcal{T}.$$

■

10. Monte Carlo Experiments

In this section we study the behavior of our test in two Monte Carlo experiments. In the first one, we provide evidence for overrejection of the exponential discounting model by the deterministic test of Browning (1989). In the second experiment, we provide evidence for the power of our testing procedure with respect to a fixed alternative.

10.1. Overrejection of Exponential Discounting for Browning's Deterministic Test

The objective of the Monte Carlo simulation exercise is to test the performance of the methodological procedure developed in this paper against the deterministic benchmark. We are going to provide evidence that a data set generated by a random exponential discounter, when contaminated with measurement error, will be erroneously rejected by deterministic methodologies at the individual level for a sizable fraction of the sample (Blow et al., 2013, Browning, 1989). However, our test will not reject it.

We choose our simulation configuration setup to match those of the single-individual household characteristics in our application. The Monte Carlo exercise will deal with a moderate size data set of $n = 200$ individuals to show that it works in a data set of the same size in our application. The sample size is $n \in \{200, 1500\}$, the time period is $\mathcal{T} = \{0, 1, 2, 3\}$, and we consider $L = 17$ goods.

We use the same discounted prices $\{\rho_{i,t}\}_{i=1}^{n}$ as the ones given in Adams et al. (2014).[34] These are the prices faced by the single-individual households in our application. We consider consumers with the constant elasticity of substitution (CES) instantaneous utility

$$u(c_t) = \sum_{l=1}^{L} \frac{c_{t,l}^{1-\sigma}}{1-\sigma},$$

where $\boldsymbol{\sigma} \sim U[1/15, 100]$ is heterogeneous across individuals. The consumers are exponential discounters with heterogeneous random discount factor $\mathbf{d} \sim U[0.8, 1]$.

Following Browning (1989), the true consumption rule for each consumer and each realized d is given by:

$$c_{t,l}^{*} = \left(\frac{1}{d^t}\rho_{t,l}\right)^{-1/\sigma},$$

for all $l = 1, \cdots, L$ and $t \in \mathcal{T}$. Measurement error is drawn from $\boldsymbol{\epsilon}_{t,l} \sim U[0.98, 1.02]$ which implies that $E[\boldsymbol{\epsilon}_{t,l}] = 1$. Then observed consumption is equal to true consumption times the multiplicative perturbation $c_{t,l} = c_{t,l}^{*}\epsilon_{t,l}$, and we define measurement error as $w_{t,l} = c_{t,l} - c_{t,l}^{*}$. Note that the implied random measurement error $\mathbf{w}_{t,l}$ has $E[\mathbf{w}_{t,l}] = 0$ by construction. The random vector $\boldsymbol{\epsilon}$ captures incorrect consumption reporting or recording, and can be as high as 1.02 times the true consumption. This means that relative measurement error is around 2 percent. Also, we observe that assumption 1 is satisfied by the proposed measurement error:

$$\mathbb{E}\left[\mathbf{d}^{-t}\boldsymbol{\rho}_t'\mathbf{w}_t\right] = 0 \quad \forall t \in \mathcal{T}.$$

This is true because $\mathbb{E}\left[\boldsymbol{\rho}_t'\mathbf{w}_t|\boldsymbol{\rho}_t, \mathbf{d}\right] = 0$, given that $\mathbb{E}\left[\mathbf{w}_{t,l}|\boldsymbol{\rho}_t, \mathbf{d}\right] = 0$ a.s.. This produces a data set of $(\boldsymbol{\rho}_{t,i}, \mathbf{c}_{t,i})_{i=1,t\in\mathcal{T}}^{i=n}$. We replicate the experiment $m = 1000$ times. The deterministic test in Browning (1989) rejects the exponential discounting model in 61.5 (62.3) percent of the cases on average across the samples for $n = 200$ ($n = 1500$), while our methodology accepts the null hypothesis that all single households are consistent with random exponential discounting (as seen in Section 10.2).

10.2. Power Analysis

We choose our simulation configuration setup to match Section 10.1 (with a sample size of $n = 1500$). However, the consumers are sophisticated quasi-hyperbolic discounters with heterogeneous random discount factor $\mathbf{d} \sim U[0.8, 1]$. The quasi-hyperbolic discounting behavior is controlled by the present bias parameter $\beta \in \{0.5, 0.6, \ldots, 1\}$, which is the same for all consumers. When $\beta = 1$, the consumers are consistent with s/ED-rationalizability and our testing procedure should accept the null hypothesis asymptotically at least with probability $1 - \alpha$. For $\beta \neq 1$ the consumers are not consistent with s/ED-rationalizability. Thus, for $\beta \neq 1$ we should reject s/ED-rationalizability asymptotically with probability 1.

[34] We use the observed price matrix and sample from it uniformly with repetition at each Monte Carlo experiment.

Following Blow et al. (2017) the consumption rule for each consumer and realized d is given by:

$$c^*_{t,l} = \left(\frac{1}{d^t}\rho_{t,l}\prod_{i=1}^{t}\frac{1}{[1-(1-\beta)\mu_i]}\right)^{-1/\sigma}, \quad l = 1, \cdots, L;\ t \in \mathcal{T},$$

where $\mu_t \in [0,1]$ for all $t \in \mathcal{T}$ captures the individual (realized) wealth effects for each realization of income level at time t. Note that $\mu_t = \sum_{l=1}^{L} p_{t,l}\frac{\partial c^*_{t,l}}{\partial a_t}$, where a_t represents the assets at time t. Since the CES utility function implies that there is at least one normal good, it follows that $\mu_t \in [0,1]$. Therefore, we generate the data set by letting $\boldsymbol{\mu}_t$ be uniformly distributed on $[0,1]$. The randomness of $\boldsymbol{\mu}_t$ captures here the differences in wealth levels across time and across consumers. The data generating process for measurement error coincides with the one presented in Section 10.1. We conduct the experiment $m = 1000$ times for each value of β.

The results are presented in figure 2. For $\beta = 1$, as expected, the rejection rate is close to 5

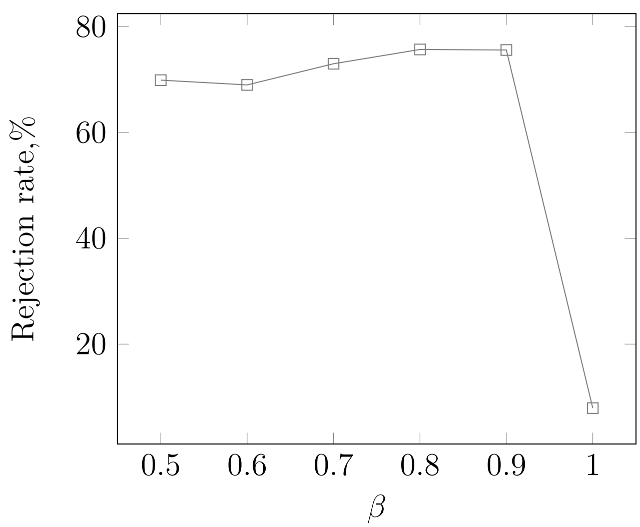

Figure 2 – Rejection rates for different levels of the present bias parameter β.

percent (7.9 percent). For $\beta \neq 1$ the rejection rates are greater or equal than 69 percent.

10.3. Pseudo-Algorithm

This pseudo-algorithm is based on Schennach's algorithm provided in GAUSS as a supplement to Schennach (2014). The actual implementation of the algorithm has been vectorized and parallelized. We indicate instances that can be vectorized, while parallelization can be done for the integration step.

1: **Step** 0

- Set Tol equal to a chosen tolerance value
- Fix $T+1$ the number of consumer experiments (let $\mathcal{T} = \{0, \cdots, T\}$).

- Fix L the number of commodities (let $\mathcal{L} = \{1, \cdots, L\}$).
- Fix $repn = (nburn, nsims)$ the number of Monte Carlo burned draws and effective draws, respectively.
- Compute the matrix $\hat{x} = (\rho_{i,t}, c_{i,t})_{i,t}$ for $i = 1, \cdots, n$ and $t = 0, \cdots, T$, where n is the sample size.
- Fix Λ the support of $(\lambda_t)_{t=0}^T$.
- Fix the bounded support $C^* = C_t^*$ for any $t \in \mathcal{T}$ (i.e., the lower bound is $0 \in \mathbb{R}^L$ and the upper bound $\overline{c}_t^* \in \mathbb{R}_+^L \setminus \{0\}$ is arbitrary large).
- Provide a particular $\eta(\cdot|\cdot) \in \mathcal{P}_{E|X}$ (i.e, the product measure of $\eta((v_t)_{t\in\mathcal{T}}, (\lambda_t)_{t\in\mathcal{T}}, (c_{l,t} - c_{l,t}^*)_{t\in\mathcal{T},l\in\mathcal{L}}|x) = f_\lambda((\lambda_t)_{t\in\mathcal{T}}) \prod_{t\in\mathcal{T}}[f_v(v_t) \prod_{l\in\mathcal{L}} f_{w|c}(c_{l,t} - c_{l,t}^*|c)]$, where the user specified density functions f_v (supported on $\mathbb{R}_+$), f_λ (supported on Λ), and $f_{w|c}$ is the measure associated with $c_{l,t} - \mathbf{c}_{l,t}^*$ where $\mathbf{c}_{l,t}^*$ has measure f_{c^*} supported on C^* and $c_{l,t}$ is given.

end Step 0.

Step 1

- Define the matrix functions: $g_I(\hat{x}_i, e) \in \mathbb{R}^k$, and $g_B(\hat{x}_i, e) \in \mathbb{R}^q$ for all $i = 1 \cdots, n$ (this step can be vectorized).
- Define the measure $\hat{\eta}(\cdot|\hat{x}_i)$, as $\hat{\eta}(e|\hat{x}_i) = \prod_{l=1}^k g_{I,l}(\hat{x}_i, e)\eta(e|\hat{x}_i)$ for all $i = 1 \cdots, n$ (this step can be vectorized).
- Set $r = -burn + 1$
- Initialize the matrix $(H_{B,i,l}(\gamma)) = 0$ for $i = 1, \cdots, n$ and $l = 1, \cdots, q$.

end Step 1.

Step 2 (Integration Step) Given $\gamma \in \mathbb{R}^q$

- For $\hat{x}_i$, draw $\hat{e}_i = ((v_{i,t})_{t\in\mathcal{T}}, (\lambda_{i,t})_{t\in\mathcal{T}}, (c_{i,l,t} - c_{i,l,t}^*)_{t\in\mathcal{T},l\in\mathcal{L}})$ proportional to $\hat{\eta}(e|\hat{x}_i)$, for all $i = 1, \cdots, n$ (this step can be vectorized). The user can draw a candidate from η efficiently and then keep only those candidates that satisfy moments g_I.
- Compute the matrix $G_B(\hat{x}, \hat{e}) = (g_B(\hat{x}_i, \hat{e}_{i,r}))_{i=1}^n \in \mathbb{R}^n \times \mathbb{R}^q$

While $r \leq nsims$

- For given $\hat{x}_i$, draw $\hat{e}_i^{jump} = ((v_{i,t})_{t\in\mathcal{T}}, (\lambda_{i,t})_{t\in\mathcal{T}}, (c_{i,l,t} - c_{i,l,t}^*)_{t\in\mathcal{T},l\in\mathcal{L}})$ proportional to $\hat{\eta}(e|\hat{x}_i)$, for all $i = 1, \cdots, n$ (this step can be vectorized).
- Compute $G_B(\hat{x}, \hat{e}^{jump}) = (g_B(\hat{x}_i, \hat{e}_i^{jump}))_{i=1}^n \in \mathbb{R}^n \times \mathbb{R}^q$
- Compute $\alpha = G_B(\hat{x}, \hat{e}^{jump})\gamma - G_B(\hat{x}, \hat{e})\gamma$

- Draw a vector $\alpha^0 = (\alpha_i^0)_{i=1}^n$ for $\alpha_i^0 \sim U[0,1]$.
- Make the set of indices $\{i_1, \cdots, i_K\} \subseteq \{1, \cdots, n\}$ such that $\alpha_{i_k}^0 < \alpha_{i_k}$.
- Set $\hat{e}_{i_k} = \hat{e}_{i_k}^{jump}$ for $i_k \in \{i_1, \cdots, i_K\}$.
- **if** $r > 0$
- Compute $H_B(\gamma) = H_B(\gamma) + (G_B(\hat{x}, \hat{e}))/nsims$
- **end if**
- Set $r = r + 1$

end While.

end Step 2**.**

Step 3 (Maximization Step)

- Define $\hat{h}_B(\gamma) = \frac{1}{n}\sum_{i=1}^n H_{B,i}(\gamma)$ using **Step** 2 for given γ.
- Define $\hat{\Omega}(\gamma) = \frac{1}{n}\sum_{i=1}^n H_{B,i}(\gamma)H_{B,i}(\gamma)' - \hat{h}_B(\gamma)\hat{h}_B(\gamma)'$.
- Find $\gamma^* = argmin_{\gamma \in \mathbb{R}^q} \hat{h}_{B,l}(\gamma)'\hat{\Omega}(\gamma)^- \hat{h}_{B,l}(\gamma)$ using NLOPT given tolerance level Tol (we use Neldermead/Bobyqa, the problem is convex in γ by construction).
- Compute $\hat{\mathrm{TS}}_n = n\hat{h}_{B,l}(\gamma^*)'\hat{\Omega}(\gamma^*)^- \hat{h}_{B,l}(\gamma^*)$.

end Step 3**.**

Sampling from $\tilde{\eta}$ can be thought as sampling from a convex polytope given a draw of $(\boldsymbol{\lambda}_t)_{t\in\mathcal{T}}$. The user can employ alternative sampling from polytopes techniques to improve speed if needed (Emiris and Fisikopoulos (2013)). We follow Schennach (2014) and use Neldermead for the optimization step. We verify the results with Bobyqa (source code supplemented by Schennach (2014)). Alternative optimization techniques that are derivative free can be used.[35]

11. Prices and Consumption Measurement Error

We consider the case when prices are also measured with error. Prices may be mismeasured because of aggregation problems. This may happen in surveys where commodities are in fact categories of goods, which implies that prices are price indexes. In particular, we will generalize the methodology of Tsur (1989) in a nonparametric fashion. We will impose a centering condition on the actual and hypothetical expenditures. Similar to the measurement error in consumption, we

[35]For our empirical application we use 7200 Monte Carlo simulations per fix value of θ_0. Using between $2000-7200$ draws we obtain similar results. These numbers are the result of trial and error and they work well in simulations.

define the price measurement error as the difference between observed prices, ρ_t, and true prices, $\boldsymbol{\rho}_t^*$:

$$\boldsymbol{\epsilon}_t = \boldsymbol{\rho}_t - \boldsymbol{\rho}_t^*,$$

where the random true prices $\boldsymbol{\rho}_t^*$ are supported on or inside $\mathbb{R}^L_{++}$ for all $t \in \mathcal{T}$. Our revealed-preference test has, as a building block, the random expenditures $\boldsymbol{\rho}_t'\mathbf{c}_s$ for all $t, s \in \mathcal{T}$. When $s = t$ we refer to this quantity as actual expenditure and when $s \neq t$ we refer to this quantity as hypothetical expenditure.

We impose the following assumption:

Assumption 6. *(Expenditure Neutrality)* $\mathbb{E}\left[\boldsymbol{\lambda}_t\boldsymbol{\rho}_t'\mathbf{c}_s\right] = \mathbb{E}\left[\boldsymbol{\lambda}_t\boldsymbol{\rho}_t^{*\prime}\mathbf{c}_s^*\right]$ *for all* $t, s \in \mathcal{T}$.

Expenditure neutrality requires that the weighted average (actual and hypothetical) expenditures are equal to the true expenditures. In terms of measurement error, Assumption 6 can be written as

$$\mathbb{E}\left[\boldsymbol{\lambda}_t(\boldsymbol{\rho}_t^{*\prime}\boldsymbol{w}_s + \boldsymbol{\epsilon}_t'\mathbf{c}_s^* + \boldsymbol{\epsilon}_t'\boldsymbol{w}_s)\right] = 0,$$

for all $s, t \in \mathcal{T}$. Thus it is sufficient to require that: (i) $\mathbb{E}\left[\boldsymbol{\lambda}_t\boldsymbol{\rho}_t^{*\prime}\boldsymbol{w}_s\right] = 0$, (ii) $\mathbb{E}\left[\boldsymbol{\lambda}_t\boldsymbol{\epsilon}_t'\mathbf{c}_s^*\right] = 0$, and (iii) $\mathbb{E}\left[\boldsymbol{\lambda}_t\boldsymbol{\epsilon}_t'\boldsymbol{w}_s\right] = 0$ for all $t, s \in \mathcal{T}$. Condition (i) corresponds to mean budget neutrality whenever prices are correctly measured and when $s = t$. Conditions (ii) and (iii) are new and can be satisfied under more primitive and natural conditions. A popular model for price measurement error is Berkson additive price measurement error:[36]

$$\boldsymbol{\rho}_t^* = \boldsymbol{\rho}_t + \boldsymbol{\eta}_t,$$

such that $\mathbb{E}\left[\boldsymbol{\eta}_t\right] = 0$. In this case we have that $\boldsymbol{\epsilon}_t = -\boldsymbol{\eta}_t$ for all $t \in \mathcal{T}$. In our framework, expenditure neutrality holds if the Berkson perturbation is independent of observed consumption and measurement error in consumption conditional on the marginal utility of wealth (i.e., $\boldsymbol{\eta}_t \perp\!\!\!\perp ((\boldsymbol{c}_t)_{t\in\mathcal{T}}, (\boldsymbol{w}_t)_{t\in\mathcal{T}})|\boldsymbol{\lambda}_t)$.

The presence of measurement error in both prices and consumption is an extreme case that requires a stronger centering conditions in the form of Condition (i) for the case of $t \neq s$. Absent expenditure neutrality, it is not clear that any revealed-preference test will have empirical bite in such a difficult environment. However, under expenditure neutrality, there is significant empirical content. The following result is a trivial consequence of Lemma 2 coupled with Assumption 6.

Lemma 3. *Under Assumption 6, if a random array* $(\boldsymbol{\rho}_t, \mathbf{c}_t)_{t\in\mathcal{T}}$ *is* s/m*-rationalizable, then there exists a positive random vector* $(\mathbf{v}_t)_{t\in\mathcal{T}}$ *and a random vector* $(\boldsymbol{\lambda}_t)_{t\in\mathcal{T}}$ *supported on or inside* Λ*, such that:*

$$\mathbb{E}\left[\mathbf{v}_t\right] - \mathbb{E}\left[\mathbf{v}_s\right] \geq \mathbb{E}\left[\boldsymbol{\lambda}_t\boldsymbol{\rho}_t'(\boldsymbol{c}_t - \mathbf{c}_s)\right] \quad \forall t, s \in \mathcal{T}.$$

This result provides a practical way to test s/m-rationalizability in the presence of both measurement error in prices and measurement error in consumption. The case of moment inequalities

[36] See for example Hjertstrand (2013), Hoderlein et al. (2013).

can also be dealt with by applying the ELVIS methodology since one can always define additional latent random slack variables that convert moment inequalities to moment equalities.[37]

12. Testing for the Collective Exponential Discounting Model in Our Framework

The important contribution of Adams et al. (2014) studies a dynamic collective consumer problem to model the behavior of couple's households. The collective model considers a case in which the household maximizes a utilitarian sum of individual utilities of each member of the couple over a vector of consumption of private and (household) public goods, given the individuals' relative power within the household (Pareto weights). In this sense, each individual member of the household is an exponential discounter but the observed consumption is a result of the collective decision making process, and may not be time-consistent. We formulate a test for the collective model using our methodology. We reject the null hypothesis of consistency of the data set with the dynamic collective model assuming that the random discount factor is supported on $[0.9, 1]$ (this support is the one used in Adams et al. (2014)). However, we fail to reject the implications of the model if we allow for substantially more heterogeneous population (the support of the random discount factor is $[0.1, 1]$). This finding is important because the collective exponential discounting model presented in Adams et al. (2014) could be considered as a potential alternative to the household exponential discounting model.

Consider a household that consists of two individuals labeled by A and B. Partition the vector of goods into publicly consumed goods indexed by H and privately consumed goods indexed by I. That is, $c_t = (c'_{t,I}, c'_{t,H})'$ and $p_t = (p'_{t,I}, p'_{t,H})'$. Let $c_{t,A}$ and $c_{t,B}$ be the consumption of the privately consumed goods of individuals A and B, respectively ($c_{t,I} = c_{t,A} + c_{t,B}$). Then the **collective household problem with exponential discounting** corresponds to the maximization of

$$V_\tau(c) = \omega_A u_A(c_{\tau,A}, c_{\tau,H}) + \omega_B u_B(c_{\tau,B}, c_{\tau,H}) + \sum_{j=1}^{T-\tau} [d_A^j \omega_A u_A(c_{\tau+j,A}, c_{\tau+j,H}) + d_B^j \omega_B u_B(c_{\tau+j,B}, c_{\tau+j,H})],$$

subject to this linear intratemporal budget constraint:

$$p'_{\tau,I} c_{\tau,I} + p'_{\tau,H} c_{\tau,H} + s_t - y_t - (1 + r_t) s_{t-1} = 0,$$

where $\omega_A, \omega_B > 0$ are Pareto weights that remain constant across time and represent the bargaining power of each household member. Individual utility functions, u_A and u_B, are assumed to be

[37] Earlier versions of this paper studied a revealed-preference test based on average revealed-preference inequalities that correspond to the case studied in this section. Using this average methodology we rejected exponential discounting for the case of couples' households and fail to reject exponential discounting for the single-individual household case. If the reader is interested in these results, they are available upon request.

continuous, locally nonsatiated and concave. The individual discount factors are similarly denoted by d_A and d_B. The rest of the elements are the same as in our main model.

The quantities $c_{t,A}, c_{t,B}$ are assumed to be unobservable to the econometrician. We observe only c_t. Adams et al. (2014) propose one solution to the collective household problem above. They assume **full efficiency** in the sense that there are personalized Lindahl prices for the publicly consumed goods $p_{t,H}$ that perfectly decentralize the above problem. The Lindahl prices are $p_{t,A} \in \mathbb{R}^{L_H}_{++}$ for household member A and the analogous $p_{t,B}$ such that $p_{t,A} + p_{t,B} = p_{t,H}$.

The existence of Lindahl prices allows us to think of members of the household as autonomous (but interlinked) exponential discounters. Under full efficiency in the collective household problem, Adams et al. (2014) established the result which is the analog of Theorem 1. Similar to the case of the single-individual household, define $\rho_{t,h} = p_{t,h}/\prod_{j=1}^{t}(1+r_j)$ for $h \in \{I, H, A, B\}$.

Theorem 7. *(Adams et al. (2014)) An array $(\rho_t, c_t)_{t\in\mathcal{T}}$ can be generated by a collective household exponential discounting model with full efficiency if and only if there exist $d_A, d_B \in (0,1]$; strictly positive vectors $(v_{t,A})_{t\in\mathcal{T}}$, $(v_{t,B})_{t\in\mathcal{T}}$; individual private consumption quantities $(c_{t,A}, c_{t,B})_{t\in\mathcal{T}}$ (with $c_{t,A} + c_{t,B} = c_{t,I}$); and personalized Lindahl prices $(p_{t,A}, p_{t,B})_{t\in\mathcal{T}}$ (with $p_{t,A} + p_{t,B} = p_{t,H}$) such that for all $s, t \in \mathcal{T}$:*

$$v_{t,A} - v_{s,A} \geq d_A^{-t}\left[\rho'_{t,I}(c_{t,A} - c_{s,A}) + \rho'_{t,A}(c_{t,H} - c_{s,H})\right],$$
$$v_{t,B} - v_{s,B} \geq d_B^{-t}\left[\rho'_{t,I}(c_{t,B} - c_{s,B}) + \rho'_{t,B}(c_{t,H} - c_{s,H})\right].$$

With this result in hand, we can establish our finding in a very straightforward manner. We let $\boldsymbol{\rho}_t$ and $\mathbf{c}_t^*$ be the random vectors of deflated prices and true consumption. Finally, we define $\mathbf{d}_A$ and $\mathbf{d}_B$ as the random discount factors for household members A and B, respectively. Also, $\mathbf{u}_A, \mathbf{u}_B$ and $\boldsymbol{\omega}_A, \boldsymbol{\omega}_B$ denote the random utility functions and random Pareto weights for each household member. We keep here the assumption about the data-generating process that we maintained for the case of s/ED-rationalizability, namely, we assume that the preferences and Pareto weights remain stable for each household after being drawn from the joint distribution of $(\mathbf{u}_A, \mathbf{d}_A, \boldsymbol{\omega}_A)$ and $(\mathbf{u}_B, \mathbf{d}_B, \boldsymbol{\omega}_B)$ at the first time period. With these preliminaries in hand, we can establish and prove a stochastic analogue to the result in Adams et al. (2014).

Theorem 8. *If a random array $(\boldsymbol{\rho}_t, \mathbf{c}_t^*)_{t\in\mathcal{T}}$ is generated by a collective household with random exponential discountings under full efficiency, then there exist random variables $\mathbf{d}_A, \mathbf{d}_B$ which are both supported on or inside $[\theta_0, 1]$, and strictly positive random vectors $(\mathbf{v}_{t,A})_{t\in\mathcal{T}}$, $(\mathbf{v}_{t,B})_{t\in\mathcal{T}}$ that satisfy*

$$\mathbf{d}_A^t(\mathbf{v}_{t,A} - \mathbf{v}_{s,A}) + \mathbf{d}_B^t(\mathbf{v}_{t,B} - \mathbf{v}_{s,B}) \geq \boldsymbol{\rho}'_t(\mathbf{c}_t^* - \mathbf{c}_s^*) \text{ a.s.} \quad \forall t, s \in \mathcal{T}.$$

Theorem 8 does not provide sufficient conditions for collective rationalizability. It must be clear that we can provide a stochastic analogue of Theorem 7, but our choice has several advantages: (i) one does not need to specify which goods are consumed privately or publicly; (ii) the inequality restrictions in Theorem 8 do not depend on the unobservable Lindahl prices and private consumption vectors, which simplifies implementation; and (iii) we can maintain Assumption 1 in a very natural form. Without loss of generality we assume that household member A is asked to report the total household consumption expenditure level (which is in fact what the econometrician observes). In that case, we replace the mean budget-neutrality condition (Assumption 1) by the analogous collective mean budget-neutrality condition.

Assumption 7. *(Collective Mean Budget Neutrality)* $\mathbb{E}\left[\mathbf{d}_A^{-t}\boldsymbol{\rho}_t'\mathbf{w}_t\right] = 0$*, for all* $t \in \mathcal{T}$*.*

Under these conditions we find that the minimal value attained by the test statistic for the collective exponential discounting model for the lower bound of the support of the random discount factors $\boldsymbol{d}_A, \mathbf{d}_B$, $\theta_0 = 0.1$, is 8.2 (the p-value is 0.085), which is below the 95-percent quantile of the χ^2_4 (9.5). Thus, we fail to reject the null hypothesis that the couples' household data set is consistent with the collective exponential discounting model under the assumptions of full efficiency, common support for preferences, and the collective mean budget constraint. However, for $\theta_0 \geq 0.12$, we reject the null hypothesis of the collective household model with exponential discounting under the efficiency assumption. In fact, the test statistic for $\theta_0 = 0.12$ is 12.1 (the p-value is 0.017). This value is well above the critical value with at least 95 percent confidence level.[38] This results says that the collective model with full efficiency is rejected if we restrict the support of the random discount factors $\mathbf{d}_A, \mathbf{d}_B$ to be on or inside $[0.12, 1]$. The failure to reject the model strongly relies on the fact that there is a positive mass of very impatient individuals. This finding is in line with Adams et al. (2014) deterministic analysis of the same sample. However, the level of heterogeneity we need to rationalize the data is significantly bigger. Once again, our results, underscore the importance of taking measurement error into account when assessing the consistency of a mismeasured data set with a model.

In order to explain the differences in the level of heterogeneity in the discount factor between our findings and Adams et al. (2014) we present the following reasons. First, we emphasize that Adams et al. (2014) report that about 3 percent of couples households are inconsistent with their deterministic test. To reach this conclusion they use $[0.9, 1]$ as a support for the random individual discount factor. Our test is done at the level of the population and thus the rejection of the collective model, in our case, could be considered as a failure of the model for this fraction of households that are time-inconsistent. Second, recall that our centering condition implies that the measurement error is non-systematic. The small fraction of the population that is inconsistent with the deterministic test could have been incompatible with the collective model when allowing for measurement error that satisfies the centering condition. In that case, our test would have rejected the null hypothesis of consistency with the collective model. Third, we must point out that our test is applied to the data set without aggregation at the level of categories or commodities, and, hence, the fraction of couples' households may be different from the reported 3 percent (in the deterministic framework). In fact, Adams et al. (2014) aggregate the different commodities of the original data set into a fewer number of categories in order to better classify them as publicly

[38]We also computed the test statistic for $\theta_0 = 0.11$ (9.36), thus the test statistic is increasing with the value of θ_0 which the expected behavior for our test statistic. Also, the test statistic values for the explored θ_0 grid for the collective model are below those of the exponential discounting model for the sample of couples' households. Again this is to be expected given that the collective model generalizes the exponential discounting model.

consumed goods or privately consumed goods. In contrast, we consider the data set with its original commodities or categories without further aggregation. Our main advantage is that we do not need to aggregate commodities because our necessary conditions do not differentiate between privately and publicly consumed goods inside the household.[39] Finally, despite these differences, we fail to reject the collective model of Adams et al. (2014) when we allow for a larger level of heterogeneity in the discount factors. This is evidence in support of time-consistent behavior for each individual inside the couples' household. The later is in line with our main empirical findings.

[39] Jerison and Jerison (1994) have provided theoretical results pointing out that commodity aggregation may lead to either accept or reject static rationality more often.