Approximate Expected Utility Rationalization

Federico Echenique Taisuke Imai Kota Saito *

July 18, 2019

Abstract

We propose a new measure of deviations from expected utility theory. For any positive number e, we give a characterization of the datasets with a rationalization that is within e (in beliefs, utility, or perceived prices) of expected utility theory. The number e can then be used as a measure of how far the data is to expected utility theory. We apply our methodology to three recent experiments. Many subjects in those experiments are consistent with utility maximization, but not with expected utility maximization. Our measure of distance to expected utility is correlated with subjects' demographic characteristics.

*Echenique: Division of the Humanities and Social Sciences, California Institute of Technology, `fede@hss.caltech.edu`. Imai: Department of Economics, LMU Munich, `taisuke.imai@econ.lmu.de`. Saito: Division of the Humanities and Social Sciences, California Institute of Technology, `saito@caltech.edu`. We are very grateful to Nicola Persico, who posed questions to us that led to some of the results in this paper, and to Dan Friedman, Yves Le Yaouanq, and Pietro Ortoleva for helpful comments. We are also grateful for the feedback provided by numerous seminar audiences. This research is supported by Grant SES1558757 from the National Science Foundation. Echenique also thanks the NSF for its support through the grant CNS-1518941. Imai also acknowledges financial support by the Deutsche Forschungsgemeinschaft through CRC TRR 190.

1 Introduction

Since its beginning, revealed preference theory has dealt with the empirical content of *general* utility maximization, but has more recently turned to the empirical content of *specific* utility theories. Mostly the focus has been on expected utility: recent theoretical work seeks to characterize the observable choice behaviors that are consistent with expected utility maximization. At the same time, a number of recent empirical revealed-preference studies use data on choices under risk and uncertainty. We seek to bridge the gap between the theoretical understanding of expected utility theory, and the machinery needed to analyze experimental data on choices under risk and uncertainty.[1]

Imagine an agent making economic decisions, choosing contingent consumption given market prices and income. A long tradition in revealed preference theory studies the consistency of such choices with utility maximization, and more recent literature has investigated consistency with expected utility theory (EU). Consistency, however, is a black or white question. The choices are either consistent with EU or they are not. Our contribution is to describe the degree to which choices are consistent with EU. We propose a *measure* of the degree of a dataset's consistency with EU.

Revealed preference theory has developed measures of consistency with general utility maximization. The most widely used measure is the Critical Cost Efficiency Index (CCEI) proposed by Afriat (1972). The basic idea in the CCEI is to fictitiously decrease an agent's budget so that fewer options are revealed preferred to a given choice. The CCEI has been widely used to analyze experimental data on choices under risk and uncertainty. See, for example, Choi et al. (2007), Ahn et al. (2014), Choi et al. (2014), Carvalho et al. (2016), and Carvalho and Silverman (2017). All of these experimental studies involve subjects making decisions under risk or uncertainty, and CCEI was proposed as a measure of consistency with general utility maximization, not EU, the most commonly-used theory to explain choices under risk or uncertainty.

Of course, there is nothing wrong with studying general utility maximization in environments with risk and uncertainty, but the data is ideally suited to studying theories of choice under uncertainty, and it should be of great interest to evaluate EU using this data. The experimental studies have used CCEI, but we shall argue (on both theoretical and empirical grounds) that CCEI is not a good measure of consistency with EU. The authors of

[1] We analyze both objective expected utility theory for choice under risk, and subjective expected utility theory for choice under uncertainty.

the empirical studies have not had the proper tools to investigate consistency with EU. The purpose of our paper is to provide such a tool.

Our main contribution is to propose a measure of how far a dataset is from being consistent with EU. The measure is different from CCEI: we explain theoretically why our measure, and not CCEI, best captures the distance of a dataset to EU theory. We also argue on empirical grounds that our measure passes "smell tests" that CCEI fails. For example, CCEI ignores the manifest violations of EU where subjects make first-order stochastically dominated choices. And CCEI does not correlate well with the property of downward-sloping demand, a property that is implied by EU maximization.[2]

In the sequel, we first explain why CCEI is not a good test of consistency with EU, and give a high-level overview of our approach. After a theoretical discussion of our measure of consistency (with objective EU discussed in Section 3 and subjective EU in Section 5), we present an empirical application using data from experiments on choices under risk (Section 4).

Our empirical application has two purposes. The first is to illustrate how our method can be applied and to argue that our measure of distance to EU is useful and sensible. The second is to offer new insights into existing data. We use data from three large-scale experiments (Choi et al., 2014; Carvalho et al., 2016; Carvalho and Silverman, 2017), each with over 1,000 subjects, that involve choices under risk. Using our methodology, the data can be used to test EU theory, not just general utility maximization.

There are two main take-away messages from our empirical application. First, the data confirms that CCEI is not a good indication of compliance with EU. Among subjects with high CCEI, who are largely consistent with utility maximization, many subjects make choices that violate monotonicity with respect to first-order stochastic dominance. Our measure detects these violations of EU, where CCEI does not. Our measure correlates well with the basic property of downward-sloping demand; CCEI does not. Moreover, the correlation between closeness to EU and demographic characteristics yields intuitive results. We find that younger subjects, those who have high cognitive abilities, and those who are working, are closer to EU behavior than older, low cognitive ability, or non-working, subjects. For some of the three experiments, we also find that highly educated, high-income, and male subjects, are closer to EU.

[2]Roughly speaking, it says that prices and quantities must be inversely related, subject to certain qualifications.

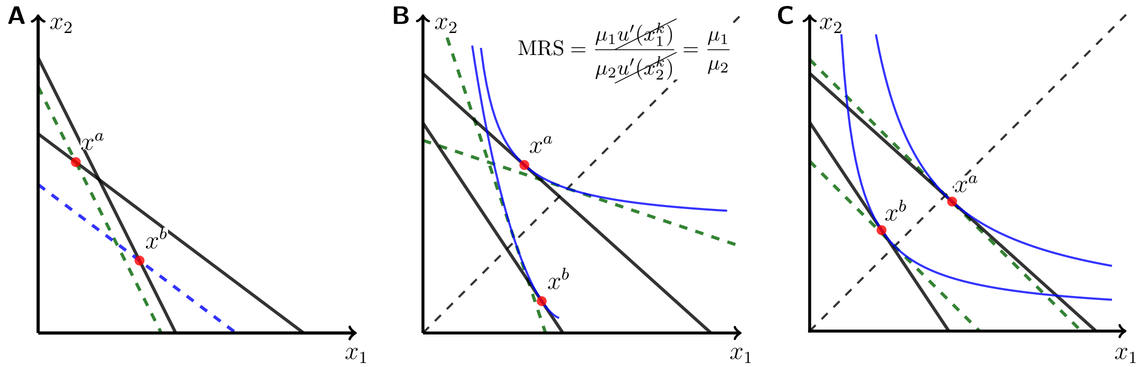

Figure 1: (A) A violation of WARP. (B) A violation of EU: $x_2^a > x_1^a$, $x_1^b > x_2^b$, and $p_1^b/p_2^b < p_1^a/p_2^a$. (C) A choice pattern consistent with EU.

1.1 How to Measure Deviations from EU

In the rest of the introduction, we lay out the theoretical arguments for why CCEI is inadequate to measure deviations from EU, and motivate our approach.

The CCEI is meant to test deviations from general utility maximization. If an agent's behavior is not consistent with utility maximization, then it cannot possibly be consistent with EU maximization. Thus it stands to reason that if an agent's behavior is far from being rationalizable as measured by CCEI, then it is also far from being rationalizable with an EU function. The problem is, of course, that an agent's behavior may be rationalizable with a general utility function but not with EU.

Broadly speaking, the CCEI proceeds by "amending" inconsistent choices through the device of changing income. This works for general utility maximization, but it is the wrong way to amend choices that are inconsistent with EU. Since EU is about getting marginal rates of substitution right, prices, not incomes, need to be changed. The problem is illustrated with a simple example in Figure 1.

Suppose that there are two states of the world, labeled 1 and 2. An agent purchases a state-contingent asset $x = (x_1, x_2)$, given Arrow-Debreu prices $p = (p_1, p_2)$ and her income. Prices and income define a budget set. In Figure 1A, we are given two choices for the agent, x^a and x^b, for two different budgets. The choices in Figure 1A are inconsistent with utility maximization: they violate the weak axiom of revealed preference (WARP). When x^b (x^a) was chosen, x^a (x^b, respectively) was strictly inside of the budget set. This violation of WARP can be resolved by shifting down the budget line associated with choice x^b to the dashed green line passing through x^a. Alternatively, the violation can be resolved by shifting down the budget line associated with choice x^a to the dashed blue line passing through x^b.

Afriat's CCEI is the smallest of the two shifts that are needed: the smallest proportion of shifting down a budget line to resolve WARP violation. Therefore, the CCEI of this dataset is given by the dashed green line passing through x^a. That is, the CCEI is $(p^b \cdot x^a)/(p^b \cdot x^b)$.

Now consider the example in Figure 1B. There are again two choices, x^a and x^b, for two different budgets. These choices do not violate WARP, and CCEI = 1 indicates perfect compliance with the theory of utility maximization. The choices in the panel are *not*, however, compatible with EU. To see why, assume that the dataset were rationalized by an expected utility: $\mu_1 u(x_1^k) + \mu_2 u(x_2^k)$, where (μ_1, μ_2) are the probabilities of the two states, and u is a (smooth) concave utility function over money. Note that the slope of a tangent line to the indifference curve at a point x^k is equal to the marginal rate of substitution (MRS): $\mu_1 u'(x_1^k)/\mu_2 u'(x_2^k)$. Moreover, at the 45-degree line (i.e., when $x_1^k = x_2^k$), the slope must be equal to $\mu_1 \cancel{u'(x_1^k)}/\mu_2 \cancel{u'(x_2^k)} = \mu_1/\mu_2$. This is a contradiction because in Figure 1B, the two tangent lines (green dashed lines) associated with x^a and x^b cross each other. Figure 1C shows an example of choices that are consistent with EU. Note that tangent lines at the 45-degree line are parallel in this case.

Importantly, the violation in Figure 1B cannot be resolved by shifting budget lines up or down, or more generally by adjusting agents' expenditures. The reason is that *the empirical content of expected utility is captured by the relation between prices and marginal rates of substitution. The slope, not the level, of the budget line, is what matters.* The basic insight comes from the equality of marginal rates of substitution and relative prices:

$$\frac{\mu_1 u'(x_1^k)}{\mu_2 u'(x_2^k)} = \frac{p_1^k}{p_2^k}. \tag{1}$$

Since marginal utility is decreasing, Equation (1) imposes a negative relation between prices and quantities. The distance to EU is directly related to how far the data is to complying with such a negative relation between prices and quantities. The formal connection is established in Theorem 2. Empirically, as we shall see, the degree of compliance of a subject's choices with this "downward sloping demand" property, goes a long way to capturing the degree of compliance of the subject's choices with EU.

Our contribution is to propose a measure of how close data is to being consistent with EU maximization. Our measure is based on the idea that marginal rates of substitution have to conform to EU maximization: whether data conform to Equation (1). If one "perturbs" marginal utility enough, then a dataset is always consistent with expected utility. Our measure is simply a measure of how large of a perturbation is needed to rationalize the data. Perturbations of marginal utility can be interpreted in three different, but equivalent, ways:

as measurement error on prices, as random shocks to marginal utility in the fashion of random utility theory (McFadden, 1974), or as perturbations to agents' beliefs. For example, if the data in Figure 1B is "e away" from being consistent with expected utility given a positive number e, then one can find beliefs μ^a and μ^b, one for each observation, so that EU is maximized for these observation-specific beliefs, and the degree of perturbation of beliefs is bounded by e.

Our measure can be applied in settings where probabilities are known and objective, for which we develop a theory in Section 3, and an application to experimental data in Section 4. It can also be applied to settings where probabilities are not known, and therefore subjective (Section 5).

Finally, we propose a statistical methodology for testing the null hypothesis of consistency with EU (Section 4.3). Our test relies on a set of auxiliary assumptions. The test indicates moderate levels of rejection of the EU hypothesis.

1.2 Related Literature

Revealed preference theory has developed tests for consistency with general utility maximization. The seminal papers include Samuelson (1938), Afriat (1967), and Varian (1982) (see Chambers and Echenique (2016) for an exposition of the basic theory).

More recent work has explored the testable implications of EU theory. This work includes Green and Srivastava (1986), Chambers et al. (2016), Kubler et al. (2014), Echenique and Saito (2015), and Polisson et al. (2017). The first four papers focus, as we do here, on rationalizability for risk-averse agents. Green and Srivastava (1986) and Chambers et al. (2016) allow for many goods in each state, which our methodology cannot accommodate. Polisson et al. (2017) present a test for EU in isolation, not jointly with risk aversion. Our assumptions are the same as in Kubler et al. (2014) and Echenique and Saito (2015).

Compared to the existing revealed preference literature on EU, our focus is different. We present a new measure of consistency with EU, not a new test. Our assumption of monetary payoffs and risk aversion is restrictive but consistent with how EU theory has been used in economics: many economic models assume risk aversion and monetary payoffs. Our results speak directly to the empirical relevance of such models. By focusing on risk aversion, we do not test EU in isolation, but the joint test of EU and risk aversion matters for many economic applications. A further motivation for focusing on risk aversion is empirical: in the data we have looked at, corner choices are very rare. This would rule out risk-seeking behavior in the context of EU. Thus, arguably, EU and risk-loving behavior would not be a serious candidate explanation of the experimental data we present in our paper.

As mentioned, the CCEI was proposed by Afriat (1972). Varian (1990) proposes a modification, and Echenique et al. (2011) and Dean and Martin (2016) propose alternative measures. Dziewulski (2018) provides a foundation for CCEI based on the model in Dziewulski (2016), which seeks to rationalize violations of utility-maximizing behavior with a model of just-noticeable differences. Compared to the literature based on the CCEI, we present an explicit model of the errors that would explain the deviation from EU. As a consequence, our measure of consistency with EU is based on a "story" for why choices are inconsistent with EU. And, as we have explained above, the nature of EU-consistent choices is poorly reflected in the CCEI's budget adjustments.

Apesteguia and Ballester (2015) propose a general method to measure the distance between theory and data in revealed preference settings. For each possible preference relation, they calculate the *swaps index*, which counts the number of alternatives that must be swapped with the chosen alternative in order for the preference relation to rationalize the data. Then, Apesteguia and Ballester (2015) consider the preference relation that minimizes the total number of swaps in all the observations, weighted by their relative occurrence in the data. Apesteguia and Ballester (2015) assume that there is a finite number of alternatives, and thus a finite number of preference relations over the set of alternatives. Because of the finiteness, they can calculate the swaps index for each preference relation and find the preference relation that minimizes the swaps index. This method by Apesteguia and Ballester (2015) is not directly applicable to our setup because in our setup, a set of alternatives is a budget set and contains infinitely many elements; moreover, the number of expected utility preferences relation is infinite. [3]

There are many other studies of revealed preference that are based on a notion of distance between the theory and the data. For example, Halevy et al. (2018) uses such distances as a guide in estimating parametric functional forms for the utility function. Polisson et al. (2017) show that a dataset is rationalizable by a model if and only if one can fit a rationalizing model to the observed data (importantly, their approach does not rely on risk aversion). They use a version of CCEI to measure deviations from the theory. Finally, de Clippel and Rozen (2019) measure consistency with general utility maximization (not EU) by way of

[3] In Appendix D.1 of Apesteguia and Ballester (2015), they consider the swaps index for expected utility preferences while assuming the finiteness of the set of alternatives. In their Appendix D.3, without axiomatization, they consider the swaps index for an infinite set of alternatives using the Lebesgue measure to "count" the number of swaps. However, they do not study the case where the number of alternatives is infinite and the preference relations are expected utility.

departures from first-order conditions, a similar approach to ours. Their work is independent and contemporaneous to ours.

2 Model

Let S be a finite set of *states*. We occasionally use S to denote the number $|S|$ of states. Let $\Delta_{++}(S) = \{\mu \in \mathbf{R}^S_{++} \mid \sum_{s=1}^S \mu_s = 1\}$ denote the set of strictly positive probability measures on S. In our model, the objects of choice are state-contingent monetary payoffs, or *monetary acts*. A monetary act is a vector in $\mathbf{R}^S_+$.

Definition 1. *A* dataset *is a finite collection of pairs* $(x, p) \in \mathbf{R}^S_+ \times \mathbf{R}^S_{++}$.

The interpretation of a dataset $(x^k, p^k)_{k=1}^K$ is that it describes K purchases of a state-contingent payoff x^k at some given vector of prices p^k, and income $p^k \cdot x^k$.

For any prices $p \in \mathbf{R}^S_{++}$ and positive number $I > 0$, the set

$$B(p, I) = \{y \in \mathbf{R}^S_+ \mid p \cdot y \leq I\}$$

is the *budget set* defined by p and I.

Expected utility theory requires a decision maker to solve the problem

$$\max_{x \in B(p,I)} \sum_{s \in S} \mu_s u(x_s), \tag{2}$$

when faced with prices $p \in \mathbf{R}^S_{++}$ and income $I > 0$, where $\mu \in \Delta_{++}(S)$ is a belief and u is a concave utility function over money. We are interested in concave u; an assumption that corresponds to risk aversion.

The belief μ will have two interpretations in our model. First, in Section 3, we shall focus on decisions taken under *risk*. The belief μ will be a known "objective" probability measure $\mu^* \in \Delta_{++}(S)$. Then, in Section 5, we study choice under *uncertainty*. Consequently, The belief μ will be a subjective beliefs, which is unobservable to us as outside observers.

When imposed on a dataset, expected utility maximization (2) may be too demanding. We are interested in situations where the model in (2) holds *approximately*. As a result, we shall relax (2) by "perturbing" some elements of the model. The exercise will be to see if a dataset is consistent with the model in which some elements have been perturbed. Specifically, we shall perturb beliefs, utilities, or prices.

First, consider a perturbation of utility u. We allow u to depend on the choice problem k and the realization of the state s. We suppose that the utility of consumption x_s in state

s is given by $\varepsilon_s^k u(x_s)$, with ε_s^k being a (multiplicative) perturbation in utility. To sum up, given price p and income I, a decision maker solves the problem

$$\max_{x \in B(p,I)} \sum_{s \in S} \mu_s \varepsilon_s^k u(x_s),$$

when faced with prices $p \in \mathbf{R}_{++}^S$ and income $I > 0$. Here $\{\varepsilon_s^k\}$ is a set of perturbations, and u is, as before, a concave utility function over money.

In the second place, consider a perturbation of beliefs. We allow μ to be different for each choice problem k. That is, given price p and income I, a decision maker solves the problem

$$\max_{x \in B(p,I)} \sum_{s \in S} \mu_s^k u(x_s), \tag{3}$$

when faced with prices $p \in \mathbf{R}_{++}^S$ and income $I > 0$, where $\{\mu^k\} \subset \Delta_{++}(S)$ is a set of beliefs and u is a concave utility function over money.

Finally, consider a perturbation of prices. Our consumer faces perturbed prices $\tilde{p}_s^k = \varepsilon_s^k p_s^k$, with a perturbation ε_s^k that depends on the choice problem k and the state s. Given price p and income I, a decision maker solves the problem

$$\max_{x \in B(\tilde{p},I)} \sum_{s \in S} \mu_s u(x_s),$$

when faced with income $I > 0$ and the perturbed prices $\tilde{p}_s^k = \varepsilon_s^k p_s^k$ for each $k \in K$ and $s \in S$.

Observe that our three sources of perturbations have different interpretations, each can be traced back to a long-standing tradition for how errors are introduced in economic models. Perturbed prices can be thought of a prices subject to measurement error, measurement error being a very common source of perturbations in econometrics (Griliches, 1986). Perturbed utility is an instance of random utility models (McFadden, 1974). Finally, perturbations of beliefs can be thought of as a kind of random utility, or as an inability to exactly use probabilities. Note that we perturb one source at a time and do not consider combinations of perturbations.

3 Perturbed Objective Expected Utility

In this section, we discuss choice under risk: there exists a known "objective" belief $\mu^* \in \Delta_{++}(S)$ that determines the realization of states. The experiments we discuss in Section 4 are all on choice under risk.

As mentioned above, we go through each of the sources of perturbation: beliefs, utility, and prices. We seek to understand how large a perturbation has to be in order to rationalize a dataset. It turns out that, for this purpose, all sources of perturbations are equivalent.

3.1 Belief Perturbation

Deviations from EU are accommodated by allowing a different belief at each observation. So we assume a belief μ^k for each choice k, and allow μ^k to differ from the objective μ^*. We seek to understand how much the belief μ^k deviates from the objective belief μ^* by evaluating how far the ratio,

$$\frac{\mu_s^k/\mu_t^k}{\mu_s^*/\mu_t^*},$$

where $s \neq t$, differs from 1. If the ratio is larger (smaller) than one, then it means that in choice k, the decision maker believes the relative likelihood of state s with respect to state t is larger (smaller, respectively) than what he should believe, given the objective belief μ^*.

Given a non-negative number e, we say that a dataset is e-belief-perturbed objective expected utility (OEU) rational, if it can be rationalized using expected utility with perturbed beliefs for which the relative likelihood ratios do not differ by more than e from their objective equivalents. Formally:

Definition 2. *Let $e \in \mathbf{R}_+$. A dataset $(x^k, p^k)_{k=1}^K$ is* e-belief-perturbed OEU rational *if there exist $\mu^k \in \Delta_{++}$ for each $k \in K$, and a concave and strictly increasing function $u : \mathbf{R}_+ \to \mathbf{R}$, such that, for all k,*

$$y \in B(p^k, p^k \cdot x^k) \implies \sum_{s \in S} \mu_s^k u(y_s) \leq \sum_{s \in S} \mu_s^k u(x_s^k).$$

and for each $k \in K$ and $s, t \in S$,

$$\frac{1}{1+e} \leq \frac{\mu_s^k/\mu_t^k}{\mu_s^*/\mu_t^*} \leq 1 + e. \tag{4}$$

When $e = 0$, e-belief-perturbed OEU rationality requires that $\mu_s^k = \mu_s^*$ for all s and k, so the case of exact consistency with expected utility is obtained with a zero bound of belief perturbations. Moreover, it is easy to see that by taking e to be large enough, any dataset can be e-belief-perturbed rationalizable.

We should note that e bounds belief perturbations for all states and observations. As such, it is sensitive to extreme observations and outliers (the CCEI is also subject to this critique: see Echenique et al., 2011). In our empirical application, we carry out a robustness analysis to account for such sensitivity (see Appendix E.2).

Finally, we mention a potential relationship with models of nonexpected utility. One could think of rank-dependent utility, for example, as a way of allowing agent's beliefs to adapt to his observed choices. However, unlike e-belief-perturbed OEU, the nonexpected utility theory requires some consistencies on the dependency. For example, for the case of rank-dependent utility, the agent's belief over the states is affected by the ranking of the outcomes across states.

3.2 Price Perturbation

We now turn to perturbed prices: think of them as prices measured with error. The perturbation is a multiplicative noise term ε_s^k to the Arrow-Debreu state price p_s^k. Thus, perturbed state prices are $\varepsilon_s^k p_s^k$. Note that if $\varepsilon_s^k = \varepsilon_t^k$ for all s, t, then introducing the noise does not affect anything because it only changes the scale of prices. In other words, what matters is how perturbations affect relative prices, that is $\varepsilon_s^k / \varepsilon_t^k$.

We can measure how much the noise ε^k perturbs relative prices by evaluating how much the ratio,

$$\frac{\varepsilon_s^k}{\varepsilon_t^k},$$

where $s \neq t$, differs from 1.

Definition 3. *Let $e \in \mathbf{R}_+$. A dataset $(x^k, p^k)_{k=1}^K$ is e-*price-perturbed OEU rational *if there exists a concave and strictly increasing function $u : \mathbf{R}_+ \to \mathbf{R}$, and $\varepsilon^k \in \mathbf{R}_+^S$ for each $k \in K$ such that, for all k,*

$$y \in B(\tilde{p}^k, \tilde{p}^k \cdot x^k) \implies \sum_{s \in S} \mu_s^* u(y_s) \leq \sum_{s \in S} \mu_s^* u(x_s^k),$$

where for each $k \in K$ and $s \in S$

$$\tilde{p}_s^k = p_s^k \varepsilon_s^k$$

and for each $k \in K$ and $s, t \in S$

$$\frac{1}{1+e} \leq \frac{\varepsilon_s^k}{\varepsilon_t^k} \leq 1 + e. \tag{5}$$

It is without loss of generality to add an additional restriction that $\tilde{p}_k \cdot x_k = p_k \cdot x_k$ for each $k \in K$ because what matters are the relative prices.

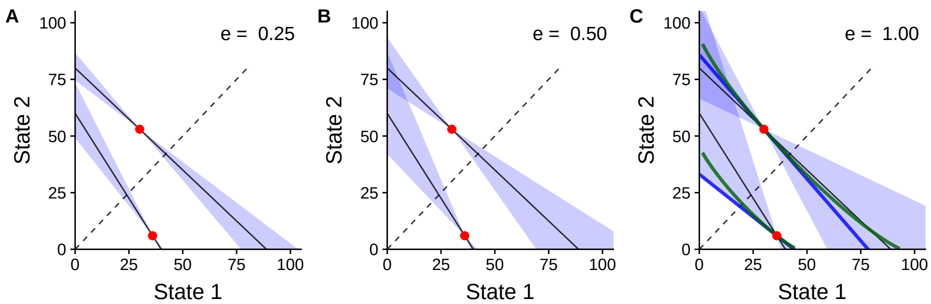

Figure 2: (A-C) Illustration of perturbed budget sets with $e \in \{0.25, 0.5, 1\}$. (C) Example of price-perturbed OEU rationalization.

The idea is illustrated in Figure 2A-D. The figure shows how the perturbations to relative prices affect budget lines, under the assumption that $|S| = 2$. For each value of $e \in \{0.25, 0.5, 1\}$ and $k \in K$, the blue area is the set $\{x \in \mathbf{R}^S_+ \mid x \cdot \tilde{p}^k = x^k \cdot \tilde{p}^k \text{ and } (5)\}$ of perturbed budget lines. The dataset in the figure is the same as in Figure 1B, which is not rationalizable with any expected utility function.

Figure 2C illustrates how we rationalize the dataset in Figure 1B. The blue bold lines are perturbed budget lines and the green bold curves are (fixed) indifference curves passing through each of the x^k in the data. The blue shaded areas are the sets of perturbed budget lines bounded by $e = 1$. Perturbed budget lines needed to rationalize the choices are indicated with blue bold lines. Since these are inside the shaded areas, the dataset is price-perturbed OEU rational with $e = 1$.

3.3 Utility Perturbation

Finally, we turn to perturbed utility. As explained above, perturbations are multiplicative and take the form $\varepsilon^k_s u(x^k_s)$. It is easy to see that this method is equivalent to belief perturbation.[4] As for price perturbations, we seek to measure how much the ε^k perturbs utilities at choice problem k by evaluating how much the ratio,

$$\frac{\varepsilon^k_s}{\varepsilon^k_t},$$

where $s \neq t$, differs from 1.

Definition 4. *Let $e \in \mathbf{R}_+$. A dataset $(x^k, p^k)^K_{k=1}$ is* e-utility-perturbed OEU rational *if there exists a concave and strictly increasing function $u : \mathbf{R}_+ \to \mathbf{R}$ and $\varepsilon^k \in \mathbf{R}^S_+$ for each $k \in K$*

[4]We consider state-contingent perturbations. As such, perturbed utilities fall outside of the domain of EU theory. We thank Jose Apesteguía and Miguel Ballester for pointing this out to us.

such that, for all k,

$$y \in B(p^k, p^k \cdot x^k) \implies \sum_{s \in S} \mu_s^* \varepsilon_s^k u(y_s) \le \sum_{s \in S} \mu_s^* \varepsilon_s^k u(x_s^k),$$

and for each $k \in K$ and $s, t \in S$

$$\frac{1}{1+e} \le \frac{\varepsilon_s^k}{\varepsilon_t^k} \le 1 + e. \tag{6}$$

3.4 Equivalence of Belief, Price, and Utility Perturbations

The first observation we make is that the three sources of perturbations are equivalent, in the sense that for any e a dataset is e-perturbed rationalizable according to one of the sources if and only if it is also rationalizable according to any of the other sources with the same e. By virtue of this result, we can interpret our measure of deviations from OEU in any of the ways we have introduced.

Theorem 1. *Let $e \in \mathbf{R}_+$, and D be a dataset. The following are equivalent:*

- *D is e-belief-perturbed OEU rational;*
- *D is e-price-perturbed OEU rational;*
- *D is e-utility-perturbed OEU rational.*

The proof appears in Section 6. In light of Theorem 1, we shall simply say that a dataset is *e-perturbed OEU rational* if it is e-belief-perturbed OEU rational, and this will be equivalent to being e-price-perturbed OEU rational, and e-utility-perturbed OEU rational.

3.5 Characterizations

We proceed to give a characterization of the dataset that are e-perturbed OEU rational. Specifically, given $e \in \mathbf{R}_+$, we propose a revealed preference axiom and prove that a dataset satisfies the axiom if and only if it is e-perturbed OEU rational.

Before we state the axiom, we need to introduce some additional notation. In the current model, where μ^* is known and objective, what matters to an expected utility maximizer is not the state price itself, but instead the *risk-neutral* price.

Definition 5. *For any dataset $(p^k, x^k)_{k=1}^K$, the* risk neutral price $\rho_s^k \in \mathbf{R}_{++}^S$ *in choice problem k at state s is defined by*

$$\rho_s^k = \frac{p_s^k}{\mu_s^*}.$$

As in Echenique and Saito (2015), the axiom we propose involves a sequence $(x_{s_i}^{k_i}, x_{s_i'}^{k_i'})_{i=1}^n$ of pairs satisfying certain conditions.

Definition 6. *A sequence of pairs* $(x_{s_i}^{k_i}, x_{s_i'}^{k_i'})_{i=1}^n$ *is called a* test sequence *if*

(i) $x_{s_i}^{k_i} > x_{s_i'}^{k_i'}$ *for all* i*;*

(ii) each k *appears as* k_i *(on the left of the pair) the same number of times it appears as* k_i' *(on the right).*

Echenique and Saito (2015) provide an axiom for OEU rationalization, termed the Strong Axiom for Revealed Objective Expected Utility (SAROEU), which states that for any test sequence $(x_{s_i}^{k_i}, x_{s_i'}^{k_i'})_{i=1}^n$, we have

$$\prod_{i=1}^{n} \frac{\rho_{s_i}^{k_i}}{\rho_{s_i'}^{k_i'}} \leq 1. \tag{7}$$

SAROEU is equivalent to the axiom provided by Kubler et al. (2014).

It is easy to see why SAROEU is necessary for OEU rationalization. Assuming (for simplicity of exposition) that u is differentiable, the first-order condition of the maximization problem (2) for choice problem k is

$$\lambda^k p_s^k = \mu_s^* u'(x_s^k), \text{ or equivalently, } \rho_s^k = \frac{u'(x_s^k)}{\lambda^k},$$

where $\lambda^k > 0$ is a Lagrange multiplier.

By substituting this equation on the left hand side of (7), we have

$$\prod_{i=1}^{n} \frac{\rho_{s_i}^{k_i}}{\rho_{s_i'}^{k_i'}} = \prod_{i=1}^{n} \frac{\lambda^{k_i'}}{\lambda^{k_i}} \cdot \prod_{i=1}^{n} \frac{u'(x_{s_i}^{k_i})}{u'(x_{s_i'}^{k_i'})} \leq 1.$$

To see that this term is smaller than 1, note that the first term of the product of the λ-ratios is equal to one because of the condition (ii) of the test sequence: all λ^k must cancel out. The second term of the product of u'-ratio is less than one because of the concavity of u, and the condition (i) of the test sequence (i.e., $u'(x_{s_i}^{k_i})/u'(x_{s_i'}^{k_i'}) \leq 1$). Thus, SAROEU is implied. It is more complicated to show that SAROEU is sufficient (see Echenique and Saito, 2015).

Now, e-perturbed OEU rationality allows the decision maker to use different beliefs $\mu^k \in \Delta_{++}(S)$ for each choice problem k. Consequently, SAROEU is not necessary for e-perturbed OEU rationality. To see that SAROEU can be violated, note that the first-order condition of the maximization (3) for choice k is as follows: there exists a positive number (Lagrange multiplier) λ^k such that for each $s \in S$,

$$\lambda^k p_s^k = \mu_s^k u'(x_s^k), \text{ or equivalently, } \rho_s^k = \frac{\mu_s^k}{\mu_s^*}\frac{u'(x_s^k)}{\lambda^k}.$$

Suppose that $x_s^k > x_t^k$. Then (x_s^k, x_t^k) is a test sequence (of length one) according to Definition 6. We have

$$\frac{\rho_s^k}{\rho_t^k} = \left(\frac{\mu_s^k}{\mu_s^*}\frac{u'(x_s^k)}{\lambda^k}\right) \Big/ \left(\frac{\mu_t^k}{\mu_t^*}\frac{u'(x_t^k)}{\lambda^k}\right) = \frac{u'(x_s^k)}{u'(x_t^k)}\frac{\mu_s^k/\mu_t^k}{\mu_s^*/\mu_t^*}.$$

Even though $x_s^k > x_t^k$ implies the first term of the ratio of u' is less than one, the second term can be strictly larger than one. When x_s^k is close enough to x_t^k, the first term is almost one while the second term can be strictly larger than one. Consequently, SAROEU can be violated.

However, by (4), we know that the second term is bounded by $1+e$. So we must have

$$\frac{\rho_s^k}{\rho_t^k} \leq 1+e.$$

In general, for a sequence $(x_{s_i}^{k_i}, x_{s_i'}^{k_i'})_{i=1}^n$ of pairs, one may suspect that the bound is calculated as $(1+e)^n$. This is not true because if x_s^k appears both as $x_{s_i}^{k_i}$ for some i (on the left of the pair) and as $x_{s_j'}^{k_j'}$ for some j (on the right of the pair), then all μ_s^k can be canceled out. What matters is the number of times x_s^k appears without being canceled out. This number can be defined as follows.

Definition 7. *Consider any sequence $(x_{s_i}^{k_i}, x_{s_i'}^{k_i'})_{i=1}^n$ of pairs. Let $(x_{s_i}^{k_i}, x_{s_i'}^{k_i'})_{i=1}^n \equiv \sigma$. For any $k \in K$ and $s \in S$,*

$$d(\sigma, k, s) = \#\{i \mid x_s^k = x_{s_i}^{k_i}\} - \#\{i \mid x_s^k = x_{s_i'}^{k_i'}\},$$

and

$$m(\sigma) = \sum_{s \in S} \sum_{k \in K: d(\sigma,k,s)>0} d(\sigma, k, s).$$

Note that, if $d(\sigma, k, s)$ is positive, then $d(\sigma, k, s)$ is the number of times μ_s^k appears as a numerator without being canceled out. If it is negative, then $d(\sigma, k, s)$ is the number of times μ_s^k appears as a denominator without being canceled out. So $m(\sigma)$ is the "net" number of terms such as μ_s^k/μ_t^k that are present in the numerator. Thus the relevant bound is $(1+e)^{m(\sigma)}$.

Given the discussion above, it is easy to see that the following axiom is necessary for e-perturbed OEU rationality.

Axiom 1 (e-Perturbed Strong Axiom for Revealed Objective Expected Utility (e-PSAROEU)). *For any test sequence of pairs* $(x^{k_i}_{s_i}, x^{k'_i}_{s'_i})^n_{i=1} \equiv \sigma$*, we have*

$$\prod_{i=1}^{n} \frac{\rho^{k_i}_{s_i}}{\rho^{k'_i}_{s'_i}} \leq (1+e)^{m(\sigma)}.$$

The main result of this section is to show that the axiom is also sufficient.

Theorem 2. *Given* $e \in \mathbf{R}_+$*, and let* D *be a dataset. The following are equivalent:*

- D *is* e*-belief-perturbed OEU rational.*
- D *satisfies* e*-PSAROEU.*

The proof appears in Section 6.

Axioms like e-PSAROEU can be interpreted as a statement about downward-sloping demand (see Echenique et al., 2016). For example $(x^k_s, x^k_{s'})$ with $x^k_s > x^k_{s'}$ is a test sequence. If risk neutral prices satisfy $\rho^k_s > \rho^k_{s'}$, then the dataset violates downward-sloping demand. Now e-PSAROEU measures the extent of the violation by controlling the size of $\rho^k_s/\rho^k_{s'}$.

In its connection to downward-sloping demand, Theorem 2 formalizes the idea of testing OEU through the correlation of risk-neutral prices and quantities: see Friedman et al. (2018) and our discussion in Section 4.2. Theorem 2 and the axiom e-PSAROEU give the precise form that the downward-sloping demand property takes in order to characterize OEU, and provide a non-parametric justification to the practice of analyzing the correlation of prices and quantities.

As mentioned, 0-PSAROEU is equivalent to SAROEU. When $e = \infty$, the e-PSAROEU always holds because $(1+e)^{m(\sigma)} = \infty$.

Given a dataset, we shall calculate the *smallest* e for which the dataset satisfies e-PSAROEU. It is easy to see that such a minimal level of e exists. [5] We explain in Appendices B and C how it is calculated in practice.

Definition 8. Minimal e*, denoted* e_**, is the smallest* $e' \geq 0$ *for which the data satisfies* e'*-PSAROEU.*

The number e_* is a crucial component of our empirical analysis. Importantly, it is the basis of a statistical procedure for testing the null hypothesis of OEU rationality.

[5] In Appendix B, we show that e_* can be obtained as a solution of minimization of a continuous function on a compact space. So the minimum exists.

As mentioned above, e_* is a bound that has to hold across all observations, and therefore may be sensitive to extreme outliers. It is, however, easy to check the sensitivity of the calculated e_* to an extreme observation. One can, for example, re-calculate e_* after dropping one or two observations, and look for large changes.

Finally, e_* depends on the prices and the objective probability which a decision maker faces. In particular, it is clear from e-PSAROEU that $1 + e$ is bounded by the maximum ratio of risk-neutral prices (i.e., $\max_{k,k' \in K, s,s' \in S} \rho_s^k / \rho_{s'}^{k'}$).

We should mention that Theorem 2 is similar in spirit to some of the results in Allen and Rehbeck (2018), who consider approximate rationalizability of quasilinear utility. They present a revealed preference characterization with a measure of error "built in" to the axiom, similar to ours, which they then use as an input to a statistical test. The two papers were developed independently, and since the models in question are very different, the results are unrelated.

4 Testing (Objective) Expected Utility

We apply our methodology to data from three large-scale online experiments. The experiments were implemented through representative surveys, and involved objective risk, not uncertainty. The data are taken from Choi et al. (2014, hereafter CKMS), Carvalho et al. (2016, hereafter CMW), and Carvalho and Silverman (2017, hereafter CS). All three experiments share a common experimental structure, the portfolio allocation task introduced by Loomes (1991) and Choi et al. (2007).

It is worth mentioning again that the three studies focus on CCEI as a measure of violation of basic rationality. We shall instead look at OEU, and use e_* as our measure of violations of OEU. Our procedure for calculating e_* is explained in Appendices B and C.

4.1 Datasets

In the experiments, subjects were presented with a sequence of decision problems under risk in a graphical illustration of a two-dimensional budget line. They were then asked to select a point (x_1, x_2), an "allocation," by clicking on the budget line (subjects were therefore forced to exhaust the income). The coordinates of the selected point represent an allocation of points between "accounts" 1 and 2. They received the points allocated to one of the accounts, determined at random with equal chance ($\mu_1^* = \mu_2^* = 0.5$). Subjects faced a total

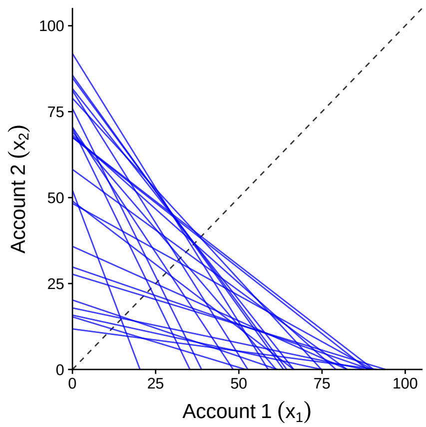

Figure 3: Sample budget lines. A set of 25 budgets from one real subject in Choi et al. (2014).

of 25 budgets, as illustrated in Figure 3.

We note some interpretations of the design that matter for our posterior discussion. First, points on the 45-degree line correspond to equal allocations between the two accounts, and therefore involve no risk. The 45-degree line is the "full insurance" line. Second, we can interpret the slope of a budget line as a price in the usual sense: if the x_2-intercept is larger than the x_1-intercept, points in the account 2 are "cheaper" than those in the account 1.

Choi et al. (2014) implemented the task using the instrument of the CentERpanel, randomly recruiting subjects from the entire panel sample in the Netherlands. Carvalho et al. (2016) administered the task using the GfK KnowledgePanel, a representative panel of the adult U.S. population. Carvalho and Silverman (2017) used the Understanding America Study panel. The number of subjects completed the task in each study is 1,182 in CKMS, 1,119 in CMW, and 1,423 in CS.

The survey instruments in these studies allowed them to collect a wide variety of individual demographic and economic information from the respondents. The main sociodemographic information they obtained include gender, age, education level, household income, occupation, and household composition.

The selection of 25 budget lines was independent across subjects in CKMS (i.e., the subjects were given different sets of budget lines), fixed in CMW (i.e., all subjects saw the same set of budgets), and semi-randomized across subjects in CS (i.e., each subject drew one of the 10 sets of 25 budgets).

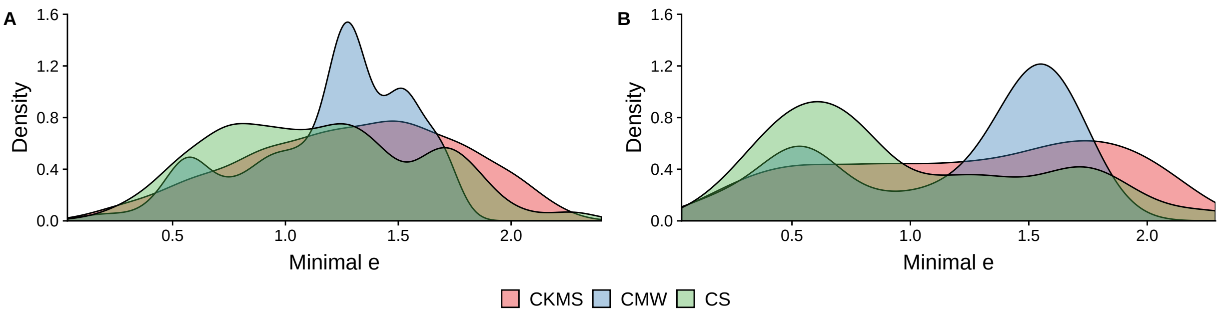

Figure 4: Kernel density estimations of e_*. (A) all subjects. (B) The subsample of subjects with CCEI = 1.

4.2 Results

Summary statistics. We exclude five subjects who are "exactly" OEU rational, leaving us a total of 3,719 subjects in the three experiments. About 76% of subjects never chose corners of the budget lines, and there are only 77 subjects (two percent of the entire sample) who chose corners in more than half of the 25 questions. Given these observations, our focus on risk aversion does not seem to be too restrictive in this environment.

We calculate e_* for each individual subject. The distributions of e_* are displayed in Figure 4A. The CKMS sample has a mean e_* of 1.289, and a median of 1.316. The CMW subjects have a mean of 1.189 and a median of 1.262, while the CS sample has a mean of 1.143 and a median of 1.128.[6]

Recall that the smaller a subject's e_* is, the closer are her choices to OEU rationality. Of course, it is hard to exactly interpret the magnitude of e_*, a problem that we turn to in Section 4.3.

Downward-sloping demand and e_*. Perturbations in beliefs, prices, or utility, seek to accommodate a dataset so that it is OEU rationalizable. The accommodation can be seen as correcting a mismatch of relative prices and marginal rates of substitution: recall our discussion in the introduction. Another way to see the accommodation is through the relation between prices and quantities. Our revealed preference axiom, e-PSAROEU, bounds certain deviations from downward-sloping demand. The minimal e is therefore a measure of the kinds of deviations from downward-sloping demand that are crucial to OEU rationality.

Figure 5 illustrates this idea. We calculate the Spearman's correlation coefficient between $\log(x_2/x_1)$ and $\log(p_2/p_1)$ for each subject in the datasets.[7] Roughly speaking, downward-

[6]Since e_* depends on the design of set(s) of budgets, comparing e_* across studies requires caution.

[7]Note that $\log(x_2/x_1)$ is not defined at the corners. We thus adjust corner choices (less than 5% of all

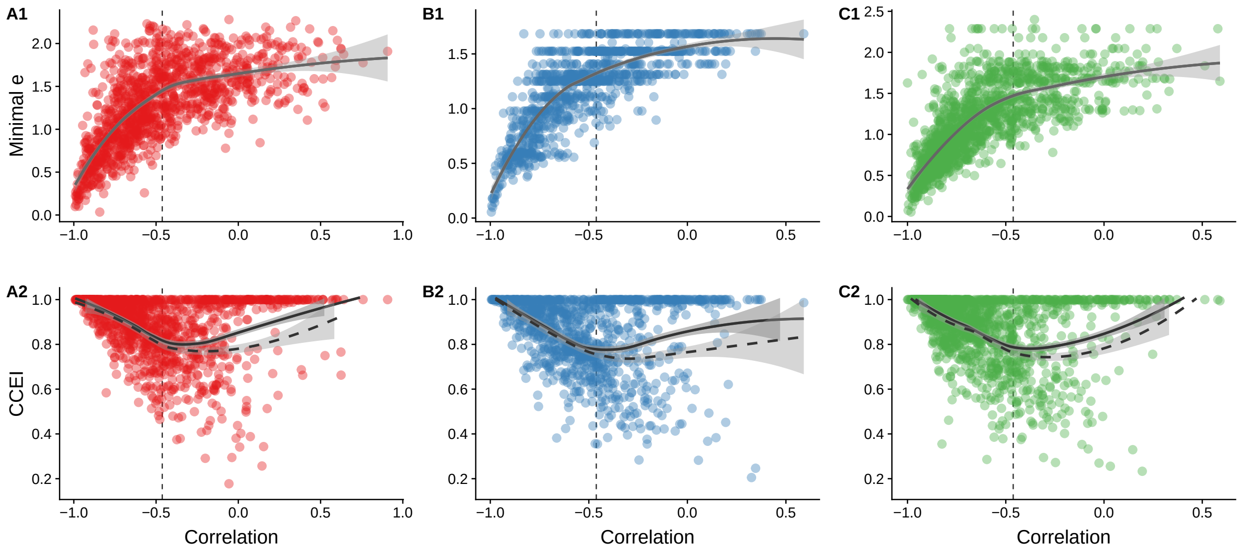

Figure 5: Correlation between $\log(x_2/x_1)$ and $\log(p_2/p_1)$ and measures of rationality: (A) CKMS, (B) CMW, (C) CS. The vertical dashed line indicates the threshold below which Spearman's correlation is significantly negative (one-sided, at the 1% level). Solid curves represent LOESS smoothing with 95% confidence bands. Dashed curves in the second row represent LOESS smoothing excluding subjects with CCEI = 1.

sloping demand corresponds to the correlation between changes in quantities $\log(x_2/x_1)$, and changes in prices $\log(p_2/p_1)$, being negative. The correlation is close to zero if subjects do not respond to price changes.

The top row of Figure 5 confirms that e_* and the correlation between prices and quantities are closely related. This means that subjects with smaller e_* tend to exhibit downward-sloping demand, while those with larger e_* are insensitive to price changes. Across all three datasets, e_* and downward-sloping demand are strongly and positively related.

The CCEI, on the other hand, is not clearly related to downward-sloping demand. As illustrated in the bottom row of Figure 5, the relation between CCEI and the correlation between prices and quantities is not monotonic. Agents who are closer to complying with utility maximization do not necessarily display a stronger negtive correlation between prices and quantities. The finding is consistent with our comment about CCEI, e_*, and OEU rationality: CCEI measures the distance from utility maximization, which is related to parallel shifts in budget lines, while e_* and OEU are about the slope of the budget lines, and about a negative relation between quantities and prices. Hence, e_* reflects the characterizing properties of OEU better than CCEI.

choices) by a small constant, 0.1% of the budget in each choice, in calculation of the correlation coefficient.

We should mention that the non-monotonic relation between CCEI and downward-sloping demand seems to be partially driven by subjects who have CCEI = 1. There are 270 (22.8%) subjects whose CCEI scores equal to one in CKMS sample, 207 (18.5%) in CMW sample, and 313 (22.0%) in CS sample, respectively. Omitting such subjects weakens the non-monotone relationship. The dotted curves in the bottom row of Figure 5 look at the relation between CCEI and the correlation coefficient excluding subjects with CCEI = 1. These curves also have non-monotonic relation, but they exhibit negative relation on a wider range of the horizontal axis, and have wider confidence bands when the correlation coefficient is positive (fewer observations).

We should also mention the practice by some authors, notably Friedman et al. (2018), to evaluate compliance with OEU by looking at the correlation between risk-neutral prices and quantities. Our e_* is clearly related to that idea, and the empirical results presented in this section can be read as a validation of the correlational approach. Friedman et al. (2018) use their approach to estimate a parametric functional form, using experimental data in which they vary objective probabilities, not just prices. [8] Our approach is non-parametric, and focused on testing OEU itself, not estimating any particular utility specification.

First-order stochastic dominance and e_*. In the experiments we consider, choosing (x_1, x_2) at prices (p_1, p_2) violates *monotonicity with respect to first-order stochastic dominance (FOSD-monotonicity)* when either (i) $p_1 > p_2$ and $x_1 > x_2$ or (ii) $p_2 > p_1$ and $x_2 > x_1$. Since the two states have the same objective probability in our datasets, choosing a greater payoff in the more expensive state violates monotonicity with respect to FOSD. Choices that violate FOSD-monotonicity are not uncommon in the data (see Table E.1 and Figure E.12 in Appendix E.3 for details).

Violations of monotonicity with respect to FOSD are related to downward-sloping demand, as they involve consuming more in the more expensive state. Importantly, the number of choices that violate FOSD-monotonicity is a good indicator of the distance to OEU. See the positive relation between the fraction of FOSD-monotonicity violations and e_* in the top row of Figure 6: subjects who frequently made choices violating FOSD-monotonicity tend to have larger e_* compared to those with fewer such violations. Note that OEU-rational choices must satisfy monotonicity with respect to FOSD. Indeed, choices made by five OEU-rational subjects ($e_* = 0$) in our data never violated FOSD-monotonicity.

[8] For the datasets we use, where probabilities are always fixed, the results we report in Figure 7 are analogous to what Friedman et al. (2018) report in their Figure 6.

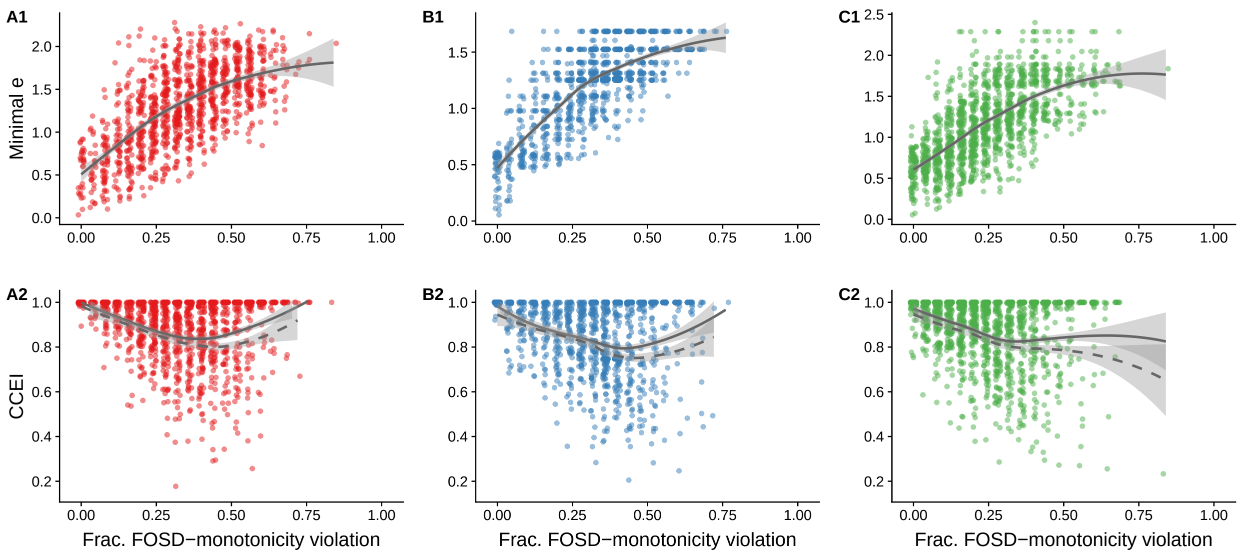

Figure 6: Violation of FOSD-monotonicity and measures of rationality. Solid curves represent LOESS smoothing with 95% confidence bands. Dashed curves in the bottom panels represent LOESS smoothing excluding subjects with CCEI = 1. Panels: (A) CKMS, (B) CMW, (C) CS.

The relation between e_* and violations of FOSD-monotonicity stands in sharp contrast with CCEI. First, choices that violate FOSD-monotonicity *can* be consistent with GARP. In fact, our data exhibits subjects that pass GARP while making choices that violate FOSD-monotonicity (an empirical fact that was first pointed out by Choi et al., 2014). The bottom row of Figure 6 shows that a substantial number of subjects with perfect compliance with GARP (CCEI = 1) make at least one violation of FOSD-monotonicity.[9] The existence of these subjects generates a nonmonotonic relationship between CCEI and the violation frequency of FOSD-monotonicity, as represented by U-shaped LOESS curves.[10]

Typical patterns of choices. We can gain some further insights into the data by considering "typical" patterns of choice. Figure 7 displays such typical patterns from selected subjects with varying degrees of e_*. Panels A-D plot observed choices from the different

[9]More than 80% of the GARP-compliant (i.e., CCEI = 1) subjects made at least one choice that violates FOSD-monotonicity (CKMS: 252/270; CMW: 173/207; CS: 265/313). Between 11 to 34% of the GARP-compliant subjects made choices violating FOSD-monotonicity in mroe than half of the 25 budgets (CKMS: 91/270; CMW: 47/207; CS: 35/313).

[10]Choi et al. (2014) propose an additional measure to make up for the problems of CCEI by combining the observed data and the *mirror-image* of the data. We consider their proposal in the appendix: see in particular, Figures E.13 and E.14 in the online appendix. The bottom line is that our conclusions continue to hold for the adjusted version of CCEI.

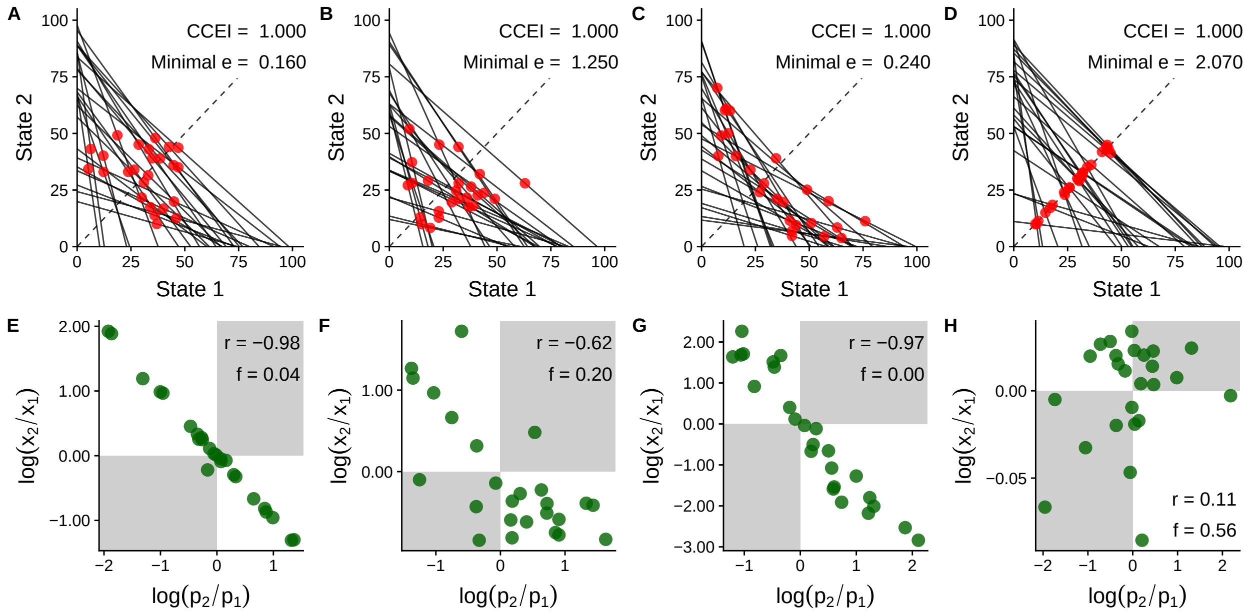

Figure 7: Choice patterns from three subjects with CCEI = 1 and varying e_*. (A-D) Observed choices. (E-H) The relation between $\log(x_2/x_1)$ and $\log(p_2/p_1)$. Choices in shaded areas violate monotonicity with respect to first-order stochastic dominance. r indicates the Pearson's correlation coefficient and f indicates the fraction of violations of FOSD-monotonicity.

budget lines, and panels E-H plot the relation between $\log(x_2/x_1)$ and $\log(p_2/p_1)$ associated with each choice pattern. The idea in the latter set of plots is that, if a subject properly responds to price changes, then as $\log(p_2/p_1)$ becomes higher, $\log(x_2/x_1)$ should become lower. This relation is also the idea in e-PSAROEU. Therefore, panels E-H in Figure 7 should have a negative slope for the subjects to be OEU rational.

Figure 7 also illustrates how e_* operates. It measures how big of an "adjustment" of prices would be needed to satisfy downward-sloping demand. Such adjustments would represent horizontal shifts in the figure. Observe that all subjects in Figure 7 have CCEI $= 1$, and are thus essentially consistent with utility maximization (ignoring the knife-edge cases of CCEI $= 1$ and inconsistency). The figure illustrates that the nature of OEU violations has little to do with CCEI.

Panel A presents a choice pattern that is "almost" consistent with OEU. The relation between $\log(x_2/x_1)$ and $\log(p_2/p_1)$ fits close to a negative line, but there is a small deviation around $\log(p_2/p_1) = 0$ which makes the subject's e_* nonzero. The choice pattern in panel B exhibits a negative slope, but with more deviations (panel F). In particular, some of the choices violate FOSD-monotonicity. Panels C and H show a pattern that does not violate

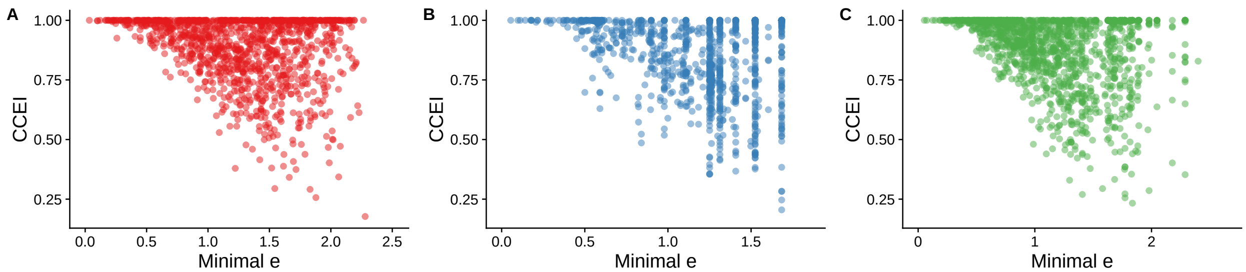

Figure 8: Correlation between e_* and CCEI from (A) CKMS, (B) CMW, and (C) CS.

FOSD-monotonicity but deviates from OEU.

The subject's choices in panel D are close to the 45-degree line. At first glance, such choices might seem to be rationalizable by a very risk-averse expected utility function. However, as panel H shows, the subject's choices deviate from downward-sloping demand, and hence cannot be rationalized by any expected utility function. One might be able to rationalize the choices made in panel D with some models of errors in choices, but not with the types of errors captured by our model.[11]

Relationship between e_* and CCEI. Comparing e_* and CCEI, we find that CCEI is not a good indication of the distance to OEU rationality. To reiterate a point we have already made, this should not be surprising as CCEI is meant to test general utility maximization, and not OEU. Nevertheless, it is interesting to see and quantify the relation between these measures in the data.

In Figure 4B, we show the distribution of e_* among subjects whose CCEI is equal to one, which varies as much as in panel A. Many subjects have CCEI equal to one, but their e_*'s are far from zero. This means that consistency with general utility maximization is not necessarily a good indication of consistency with OEU.

That said, the measures are clearly correlated. Figure 8 plots the relation between CCEI and e_*. As we expect from their definitions (*larger* CCEI and *smaller* e_* correspond to higher consistency), there is a negative and significant relation between them (Spearman's correlation coefficient: $r = -0.18$ for CKMS, $r = -0.11$ for CMW, $r = -0.35$ for CS, all

[11] This is, in our opinion, a strength of our approach. We do not ex-post seek to invent a model of errors that might rescue EU. Instead we have written down what we think are natural sources of errors and perturbation (random utility, beliefs, and measurement errors). Our results deal with what can be rationalized when these sources of errors, and only those, are used to explain the data. A general enough model of errors will, of course, render the theory untestable.

$p < 0.001$).

Notice that the variability of the CCEI scores widens as the e_* becomes larger. Obviously, subjects with a small e_* are close to being consistent with general utility maximization, and therefore have a CCEI that is close to one. However, subjects with large e_* seem to have disperse values of CCEI.

Correlation with sociodemographic characteristics. We investigate the correlation between our measure of consistency with expected utility, e_*, and various demographic variables available in the data. The exercise is analogous to CKMS's findings using CCEI.

We find that younger subjects, those who have high cognitive abilities, and those who are working, are closer to being consistent with OEU than older, low ability, or non-working, subjects. For some of the three experiments we also find that highly educated, high-income subjects, and males, are closer to OEU. Figure 9 summarizes the mean e_* along with 95% confidence intervals across several socioeconomic categories. [12] We use the same categorization as in Choi et al. (2014) to compare our results with their Figure 3.

We observe statistically significant (at the 5% level) gender differences in CMW ($t(1114) = -2.20$, $p = 0.028$) and CS ($t(1418) = -4.46$, $p < 0.001$), but not in CKMS ($t(1180) = -0.87$, $p = 0.384$). Male subjects were on average closer to OEU rationality than female subjects in the CMW and CS samples (panel A).

We find significant age effects as well. Panel B shows that younger subjects are on average closer to OEU rationality than older subjects (the comparison between age groups 16-34 and 65+ reveals statistically significant difference in all three datasets; all two-sample t-tests give $p < 0.001$).

We observe weak effects of education on e_* (panel C).[13] Subjects with higher education are on average closer to OEU than those with lower education in CKMS ($t(829) = 4.20$, $p < 0.001$), but the difference is not significant in the CMW and CS ($t(374) = 1.68$, $p = 0.094$ in CMW; $t(739) = 1.41$, $p = 0.1596$ in CS).

Panel D shows that subjects who were working at the time of the survey are on average closer to OEU than those who were not ($t(1180) = 2.24$, $p = 0.025$ in CKMS; $t(1114) = 2.43$, $p = 0.015$ in CMW; $t(1419) = 3.35$, $p = 0.001$ in CS).

[12] Figure E.16 in Appendix shows correlation between CCEI and demographic variables.

[13] The low, medium, and high education levels correspond to primary or prevocational secondary education, pre-university secondary education or senior vocational training, and vocational college or university education, respectively.

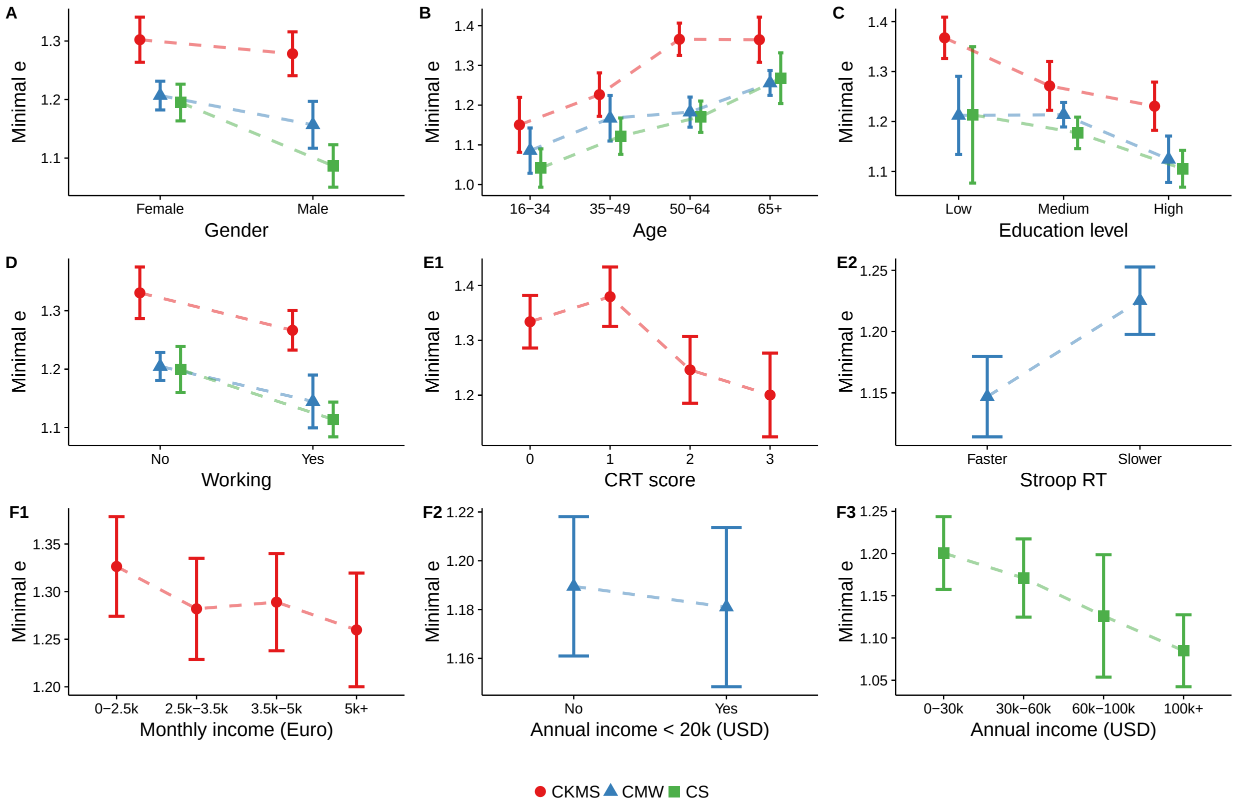

Figure 9: e_* and demographic variables.

In panels E1 and E2, we classify subjects according to their Cognitive Reflection Test score (CRT; Frederick, 2005) or average log reaction times in numerical Stroop task. [14] The average e_* for those who correctly answered two questions or more of the CRT is lower than the average for those who answered at most one question. Subjects with lower response times in the numerical Stroop task have significantly lower e_* ($t(1114) = -3.35$, $p < 0.001$).

One of the key findings in Choi et al. (2014) is that consistency with utility maximization as measured by CCEI correlates with household wealth. When we look at the relation between e_* and household income, there is a negative trend but the differences across income brackets are not statistically significant (bracket "0-2.5k" vs. "5k+", $t(533) = 1.65$, $p = 0.099$; panel F1). Panel F2 presents similar non-significance between subjects who earned

[14] CRT consists of three questions, all of which have an intuitive and spontaneous, but incorrect, answers, and a deliberative and correct answer. In the numerical Stroop task, subjects are presented with a number, such as 888, and are asked to identify the number of times the digit is repeated (in this example the answer is "3", while an "intuitive" response is "8"). It has been shown that response times in this task capture the subject's cognitive control ability.

more than 20 thousand USD annually or not in CMW sample ($t(1114) = -0.23$, $p = 0.818$). When we compare poor households (annual income less than 20 thousand USD) and wealthy households (annual income more than 100 thousand USD) from the CS sample, average e_* is significantly smaller for the latter sample ($t(887) = -3.57$, $p < 0.001$).

Robustness of the results. As discussed, e_* can be sensitive to extreme choices because it bounds perturbations for all states and observations. In a first robustness check, we recalculate e_* using subsets of observed choices after dropping one or two critical mistakes. More precisely, for each subject, we calculate e_* for all combinations of $25 - m$ ($m = 1, 2$) choices and pick the smallest e_* among them. In a second robustness check, instead of bounding all states we we calculate "average" perturbation necessary to rationalize the data to mitigate the effect of extreme mistakes. These alternative ways of calculating e_* do not change the general pattern of correlation between e_* and CCEI or e_* and demographic variables. See Appendix E.2 for details.

4.3 Minimum Perturbation Test

Our discussion so far has sidestepped one issue. How are we to interpret the absolute magnitude of e_*? When can we say that e_* is large enough to reject consistency with OEU rationality? To answer this question, we present a statistical test of the hypothesis that an agent is OEU rational. The test needs some assumptions, but it gives us a threshold level (a critical value) for e_*. Any value of e_* that exceeds the threshold indicates inconsistency with OEU at some given statistical significance level.

Our approach follows, roughly, the methodology laid out in Echenique et al. (2011) and Echenique et al. (2016). First, we adopt the price perturbation interpretation of e in Section 3.2. The advantage of doing so is that we can use the observed variability in price to get a handle on the assumptions we need to make on perturbed prices. To this end, let $D_{\text{true}} = (p^k, x^k)_{k=1}^K$ denote a dataset and $D_{\text{pert}} = (\tilde{p}^k, x^k)_{k=1}^K$ denote an "perturbed" dataset, where $\tilde{p}_s^k = p_s^k \varepsilon_s^k$ for all $s \in S$ and $k \in K$ and $\varepsilon_s^k > 0$ is a random variable. Prices $\tilde{p}^k$ are prices p^k measured with error, or misperceived.

If the *variance* of ε is large, it will be easy to accommodate a dataset as OEU rational. The larger is the variance of ε, the larger the magnitudes of e that can be rationalized as consistent with OEU. So our procedure is sensitive to the assumptions we make about the variance of ε.

To get a handle on the variance of ε, our approach is to think of an agent who mistakes

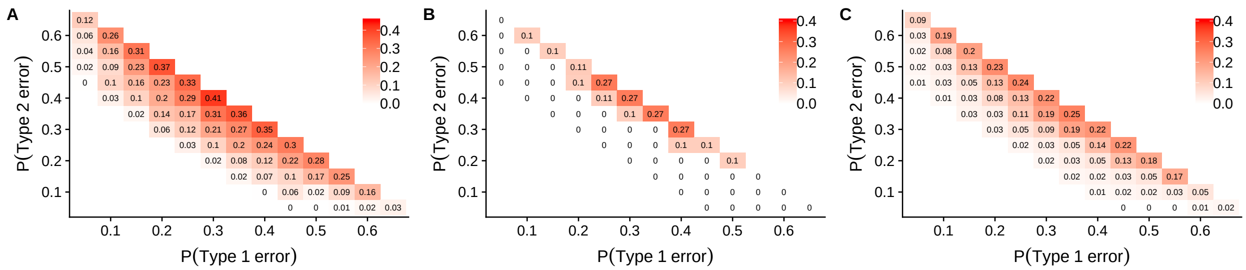

Figure 10: Rejection rates under each combination of type I and type II error probabilities (η^I, η^{II}), from CKMS sample (A), CMW sample (B), and CS sample (C).

true prices p with perturbed prices $\tilde{p}$. If the variance of ε is too large, the agent should not mistake the distribution of p and $\tilde{p}$. In other words, the distributions of p and $\tilde{p}$ should be similar enough that an agent might plausibly confuse the two. To make this operational, we imagine an agent who conducts a statistical test for the variance of prices. If the true variance of p is σ_0^2 and the implied variance of $\tilde{p}$ is $\sigma_1^2 > \sigma_0^2$, then the agent would conduct a test for the null of $\sigma^2 = \sigma_0^2$ against the alternative of $\sigma^2 = \sigma_1^2$. We want the variances to be close enough that the agent might reasonably get inconclusive results from such a test. *Specifically, we assume the sum of type I and type II errors in this test is relatively large.*[15]

The details of how we design our test are presented in Appendix D, but we can advance the main results. See Figure 10. Each panel corresponds to our results for each of the datasets. The probability of a type I error is η^I. The probability of a type II error is η^{II}. Recall that we focus on situations when $\eta^I + \eta^{II}$ is relatively large, as we want our consumer to plausibly mistake the distributions of p and $\tilde{p}$. Consider, for example, our results for CKMS. The outermost numbers assume that $\eta^I + \eta^{II} = 0.7$. For such numbers, the rejection rates range from 3% to 41%. For the CS dataset, if we look at the second line of numbers, where $\eta^I + \eta^{II} = 0.65$, we see that rejection rates range from 1% to 19%.

Overall, it is fair to say that rejection rates are modest. Smaller values of $\eta^I + \eta^{II}$ correspond to larger variances of ε, and therefore smaller rejection rates. The figure also illustrates that the conclusions of the test are very sensitive to what one assumes about variances, through the assumptions about η^I and η^{II}. But if we look at the largest rejection rates, for the largest values of $\eta^I + \eta^{II}$, we get 25% for CS, 27% for CMW, and 41% for CKMS. Hence, while many subjects in the CS, CMW and CKMS experiments are inconsis-

[15]The problem of variance is pervasive in statistical implementations of revealed preference tests, see Varian (1990), Echenique et al. (2011), and Echenique et al. (2016) for example. The use of the sum of type I and type II errors to calibrate a variance, is new to the present paper.

tent with OEU, our statistical tests would attribute such inconsistency to a mistake but at least according to our statistical test, for most subjects the rejections could be attributed to mistakes.

5 Perturbed Subjective Expected Utility

We now turn to the model of subjective expected utility (SEU), in which beliefs are not known. Instead, beliefs are subjective and unobservable. The analysis will be analogous to what we did for OEU, and therefore proceed at a faster pace. In particular, all the definitions and results parallel those of the section on OEU. The proof of the main result (the axiomatic characterization) is substantially more challenging here because both beliefs and utilities are unknown: there is a classical problem in disentangling beliefs from utility. The technique for solving this problem was introduced in Echenique and Saito (2015). The proofs of the theorems are in Appendix A.

Definition 9. *Let $e \in \mathbf{R}_+$. A dataset $(x^k, p^k)_{k=1}^K$ is* e-belief-perturbed SEU rational *if there exist $\mu^k \in \Delta_{++}$ for each $k \in K$ and a concave and strictly increasing function $u : \mathbf{R}_+ \to \mathbf{R}$ such that, for all k,*

$$y \in B(p^k, p^k \cdot x^k) \implies \sum_{s \in S} \mu_s^k u(y_s) \leq \sum_{s \in S} \mu_s^k u(x_s^k)$$

and for each $k, l \in K$ and $s, t \in S$

$$\frac{\mu_s^k / \mu_t^k}{\mu_s^l / \mu_t^l} \leq 1 + e. \tag{8}$$

Note that the definition of e-belief-perturbed SEU rationality differs from the definition of e-belief-perturbed OEU rationality, only in condition (8), establishing bounds on perturbations. Here there is no objective probability from which we can evaluate the deviation of the set $\{\mu^k\}$ of beliefs. Thus we evaluate perturbations *among* beliefs, as in (8).

Remark 1. *The constraint on the perturbation applies for each $k, l \in K$ and $s, t \in S$, so it implies for each $k, l \in K$ and $s, t \in S$*

$$\frac{1}{1+e} \leq \frac{\mu_s^k / \mu_t^k}{\mu_s^l / \mu_t^l} \leq 1 + e.$$

Hence, when $e = 0$, it must be that $\mu_s^k / \mu_t^k = \mu_s^l / \mu_t^l$. This implies that $\mu^k = \mu^l$ for a dataset that is 0*-belief perturbed SEU rational.*

Next, we propose perturbed SEU rationality with respect to prices.

Definition 10. *Let $e \in \mathbf{R}_+$. A dataset $(x^k, p^k)_{k=1}^K$ is* e-price-perturbed SEU rational *if there exist $\mu \in \Delta_{++}$ and a concave and strictly increasing function $u : \mathbf{R}_+ \to \mathbf{R}$ and $\varepsilon^k \in \mathbf{R}_+^S$ for each $k \in K$ such that, for all k,*

$$y \in B(\tilde{p}^k, \tilde{p}^k \cdot x^k) \implies \sum_{s \in S} \mu_s u(y_s) \leq \sum_{s \in S} \mu_s u(x_s^k),$$

where for each $k \in K$ and $s \in S$

$$\tilde{p}_s^k = p_s^k \varepsilon_s^k,$$

and for each $k, l \in K$ and $s, t \in S$

$$\frac{\varepsilon_s^k / \varepsilon_t^k}{\varepsilon_s^l / \varepsilon_t^l} \leq 1 + e. \tag{9}$$

Again, the definition differs from the corresponding definition of price-perturbed OEU rationality only in condition (9), establishing bounds on perturbations. In condition (9), we measure the size of the perturbations by

$$\frac{\varepsilon_s^k / \varepsilon_t^k}{\varepsilon_s^l / \varepsilon_t^l},$$

not $\varepsilon_s^k / \varepsilon_t^k$ as in (5). This change is necessary to accommodate the existence of subjective beliefs. By choosing subjective beliefs appropriately, one can neutralize the perturbation in prices if $\varepsilon_s^k / \varepsilon_t^k = \varepsilon_s^l / \varepsilon_t^l$ for all $k, l \in K$. That is, as long as $\varepsilon_s^k / \varepsilon_t^k = \varepsilon_s^l / \varepsilon_t^l$ for all $k, l \in K$, if we can rationalize the dataset by introducing the noise with some subjective belief μ, then without using the noise, we can rationalize the dataset with another subjective belief μ' such that $\varepsilon_s^k \mu_s' / \varepsilon_t^k \mu_t' = \mu_s / \mu_t$.

Finally, we define utility-perturbed SEU rationality.

Definition 11. *Let $e \in \mathbf{R}_+$. A dataset $(x^k, p^k)_{k=1}^K$ is* e-utility-perturbed SEU rational *if there exist $\mu \in \Delta_{++}$, a concave and strictly increasing function $u : \mathbf{R}_+ \to \mathbf{R}$, and $\varepsilon^k \in \mathbf{R}_+^S$ for each $k \in K$ such that, for all k,*

$$y \in B(p^k, p^k \cdot x^k) \implies \sum_{s \in S} \mu_s \varepsilon_s^k u(y_s) \leq \sum_{s \in S} \mu_s \varepsilon_s^k u(x_s^k),$$

and for each $k, l \in K$ and $s, t \in S$

$$\frac{\varepsilon_s^k / \varepsilon_t^k}{\varepsilon_s^l / \varepsilon_t^l} \leq 1 + e.$$

As in the previous section, given e, we can show that these three concepts of rationality are equivalent.

Theorem 3. *Let $e \in \mathbf{R}_+$ and D be a dataset. The following are equivalent:*

- *D is e-belief-perturbed SEU rational;*
- *D is e-price-perturbed SEU rational;*
- *D is e-utility-perturbed SEU rational.*

In light of Theorem 3, we shall speak simply of e-perturbed SEU rationality to refer to any of the above notions of perturbed SEU rationality.

Echenique and Saito (2015) prove that a dataset is SEU rational if and only if it satisfies a revealed-preference axiom termed the Strong Axiom for Revealed Subjective Expected Utility (SARSEU). SARSEU states that, for any test sequence $(x_{s_i}^{k_i}, x_{s'_i}^{k'_i})_{i=1}^n$, if each s appears as s_i (on the left of the pair) the same number of times it appears as s'_i (on the right), then

$$\prod_{i=1}^{n} \frac{p_{s_i}^{k_i}}{p_{s'_i}^{k'_i}} \leq 1.$$

SARSEU is no longer necessary for perturbed SEU-rationality. This is easy to see, as we allow the decision maker to have a different belief μ^k for each choice k, and reason as in our discussion of SAROEU. Analogous to our analysis of OEU, we introduce a perturbed version of SARSEU to capture perturbed SEU rationality. Let $e \in \mathbf{R}_+$.

Axiom 2 (e-Perturbed SARSEU (e-PSARSEU))**.** *For any test sequence $(x_{s_i}^{k_i}, x_{s'_i}^{k'_i})_{i=1}^n \equiv \sigma$, if each s appears as s_i (on the left of the pair) the same number of times it appears as s'_i (on the right), then*

$$\prod_{i=1}^{n} \frac{p_{s_i}^{k_i}}{p_{s'_i}^{k'_i}} \leq (1+e)^{m(\sigma)}.$$

We can easily see the necessity of e-PSARSEU by reasoning from the first-order conditions, as in our discussion of e-PSAROEU. The main result of this section shows that e-PSARSEU is not only necessary for e-perturbed SEU rationality, but also sufficient.

Theorem 4. *Let $e \in \mathbf{R}_+$ and D be a dataset. The following are equivalent:*

- *D is e-perturbed SEU rational;*

- *D satisfies e-PSARSEU.*

It is easy to see that 0-PSARSEU is equivalent to SARSEU, and that by choosing e to be arbitrarily large it is possible to rationalize any dataset. As a consequence, we shall be interested in finding a minimal value of e that rationalizes a dataset: such "minimal e" is also denoted by e_*.

6 Proofs of Theorems 1 and 2

6.1 Proof of Theorem 1

First we prove a lemma that implies Theorem 1, and is useful for the sufficiency part of Theorem 2. The lemma provides "Afriat inequalities" for the problem at hand.

Lemma 1. *Given $e \in \mathbf{R}_+$, and let $(x^k, p^k)_{k=1}^K$ be a dataset. The following statements are equivalent.*

(a) $(x^k, p^k)_{k=1}^K$ is e-belief-perturbed OEU rational.

(b) There are strictly positive numbers v_s^k, λ^k, μ_s^k, for $s \in S$ and $k \in K$, such that

$$\mu_s^k v_s^k = \lambda^k p_s^k, \quad \text{and} \quad x_s^k > x_{s'}^{k'} \implies v_s^k \leq v_{s'}^{k'}, \tag{10}$$

and for all $k \in K$ and $s, t \in S$,

$$\frac{1}{1+e} \leq \frac{\mu_s^k / \mu_t^k}{\mu_s^* / \mu_t^*} \leq 1 + e. \tag{11}$$

(c) $(x^k, p^k)_{k=1}^K$ is e-price-perturbed OEU rational.

(d) There are strictly positive numbers $\hat{v}_s^k$, $\hat{\lambda}^k$, and ε_s^k for $s \in S$ and $k \in K$, such that

$$\mu_s^* \hat{v}_s^k = \hat{\lambda}^k \varepsilon_s^k p_s^k, \quad \text{and} \quad x_s^k > x_{s'}^{k'} \implies \hat{v}_s^k \leq \hat{v}_{s'}^{k'},$$

and for all $k \in K$ and $s, t \in S$, $\frac{1}{1+e} \leq \frac{\varepsilon_s^k}{\varepsilon_t^k} \leq 1 + e$.

(e) $(x^k, p^k)_{k=1}^K$ is e-utility-perturbed OEU rational.

(f) There are strictly positive numbers $\hat{v}^k_s$, $\hat{\lambda}^k$, and $\hat{\varepsilon}^k_s$ for $s \in S$ and $k \in K$, such that

$$\mu^*_s \hat{\varepsilon}^k_s \hat{v}^k_s = \hat{\lambda}^k p^k_s, \quad \textit{and} \quad x^k_s > x^{k'}_{s'} \implies \hat{v}^k_s \leq \hat{v}^{k'}_{s'},$$

and for all $k \in K$ and $s, t \in S$,

$$\frac{1}{1+e} \leq \frac{\hat{\varepsilon}^k_s}{\hat{\varepsilon}^k_t} \leq 1 + e.$$

Proof. The equivalence between (a) and (b), the equivalence between (c) and (d), and the equivalence between (e) and (f) follow from arguments in Echenique and Saito (2015). The equivalence between (d) and (f) with $\varepsilon^k_s = 1/\hat{\varepsilon}^k_s$ for each $k \in K$ and $s \in S$ is straightforward. Thus, to show the result, it suffices to show that (b) and (d) are equivalent.

To show that (d) implies (b), define $v = \hat{v}$ and $\mu^k_s = \frac{\mu^*_s}{\varepsilon^k_s} / \left(\sum_{s \in S} \frac{\mu^*_s}{\varepsilon^k_s} \right)$ for each $k \in K$ and $s \in S$ and $\lambda^k = \hat{\lambda}^k / \left(\sum_{s \in S} \frac{\mu^*_s}{\varepsilon^k_s} \right)$ for each $k \in K$. Then, $\mu^k \in \Delta_{++}(S)$. Since $\mu^*_s \hat{v}^k_s = \hat{\lambda}^k \varepsilon^k_s p^k_s$, we have $\mu^k_s v^k_s = \lambda^k p^k_s$. Moreover, for each $k \in K$ and $s, t \in S$, $\frac{\varepsilon^k_s}{\varepsilon^k_t} = \frac{\mu^k_s / \mu^k_t}{\mu^*_s / \mu^*_t}$. Hence, $\frac{1}{1+e} \leq \frac{\varepsilon^k_s}{\varepsilon^k_t} \leq 1 + e$.

To show that (b) implies (d), for all $s \in S$ define $\hat{v} = v$ and for all $k \in K$, $\hat{\lambda}^k = \lambda^k$. For all $k \in K$ and $s \in S$, define $\varepsilon^k_s = \frac{\mu^*_s}{\mu^k_s}$. For each $k \in K$ and $s \in S$, since $\mu^k_s u^k_s = \lambda^k p^k_s$, we have $\mu^*_s v^k_s = \hat{\lambda}^k \varepsilon^k_s p^k_s$. Finally, for each $k \in K$ and $s, t \in S$, $\frac{\varepsilon^k_s}{\varepsilon^k_t} = \frac{\mu^*_s / \mu^k_s}{\mu^*_t / \mu^k_t} = \frac{\mu^k_t / \mu^k_s}{\mu^*_t / \mu^*_s}$. Therefore, we obtain $\frac{1}{1+e} \leq \frac{\varepsilon^k_s}{\varepsilon^k_t} \leq 1 + e$. □

6.2 Proof of the necessity direction of Theorem 2

Lemma 2. *Given $e \in \mathbf{R}_+$, if a dataset is e-belief-perturbed OEU rational, then the dataset satisfies e-PSAROEU.*

Proof. Fix any sequence $(x^{k_i}_{s_i}, x^{k'_i}_{s'_i})^n_{i=1} \equiv \sigma$ of pairs that satisfies conditions (i) and (ii) in Definition 6. By Lemma 1, there exist $v^{k_i}_{s_i}, v^{k'_i}_{s'_i}, \lambda^{k_i}, \lambda^{k'_i}, \mu^{k_i}_{s_i}, \mu^{k'_i}_{s'_i}$ such that $v^{k'_i}_{s'_i} \geq v^{k_i}_{s_i}$ and $v^{k_i}_{s_i} = \frac{\mu^*_{s_i}}{\mu^{k_i}_{s_i}} \lambda^{k_i} \rho^{k_i}_{s_i}$, and $v^{k'_i}_{s'_i} = \frac{\mu^*_{s'_i}}{\mu^{k'_i}_{s'_i}} \lambda^{k'_i} \rho^{k'_i}_{s'_i}$. Thus, we have $1 \geq \prod^n_{i=1} \frac{\lambda^{k_i} (\mu^{k'_i}_{s'_i} / \mu^*_{s'_i}) \rho^{k_i}_{s_i}}{\lambda^{k'_i} (\mu^{k_i}_{s_i} // \mu^*_{s_i}) \rho^{k'_i}_{s'_i}} = \prod^n_{i=1} \frac{\mu^{k'_i}_{s'_i} / \mu^*_{s'_i}}{\mu^{k_i}_{s_i} / \mu^*_{s_i}} \prod^n_{i=1} \frac{\rho^{k_i}_{s_i}}{\rho^{k'_i}_{s'_i}}$, where the second equality holds by condition (ii). Hence, $\prod^n_{i=1} \frac{\rho^{k_i}_{s_i}}{\rho^{k'_i}_{s'_i}} \leq \prod^n_{i=1} \frac{\mu^{k_i}_{s_i} / \mu^*_{s_i}}{\mu^{k'_i}_{s'_i} / \mu^*_{s'_i}}$.

In the following, we evaluate the right hand side. For each (k, s), we first cancel out all the terms μ^k_s that can be canceled out. Then, the number of μ^k_s's that remain in the numerator is $d(\sigma, k, s)$, as in Definition 7. Since the number of terms in the numerator and the denominator must be the same, the number of remaining fractions is

$m(\sigma) \equiv \sum_{s\in S}\sum_{k\in K:d(\sigma,k,s)>0} d(\sigma,k,s)$. So by relabeling the index i to j if necessary, we obtain $\prod_{i=1}^{n} \frac{\mu^{k_i}_{s_i}/\mu^*_{s_i}}{\mu^{k'_i}_{s'_i}/\mu^*_{s'_i}} = \prod_{j=1}^{m(\sigma)} \frac{\mu^{k_j}_{s_j}/\mu^*_{s_j}}{\mu^{k'_j}_{s'_j}/\mu^*_{s'_j}}$.

Consider the corresponding sequence $(x^{k_j}_{s_j}, x^{k'_j}_{s'_j})_{j=1}^{m(\sigma)}$. Since the sequence is obtained by canceling out x^k_s from the first element and the second element of the pairs, and since the original sequence $(x^{k_i}_{s_i}, x^{k'_i}_{s'_i})_{i=1}^{n}$ satisfies condition (ii), it follows that $(x^{k_j}_{s_j}, x^{k'_j}_{s'_j})_{j=1}^{m(\sigma)}$ satisfies condition (ii).

By condition (ii), we can assume without loss of generality that $k_j = k'_j$ for each j. Therefore, by the condition on the perturbation, $\prod_{j=1}^{m(\sigma)} \frac{\mu^{k_j}_{s_j}/\mu^*_{s_j}}{\mu^{k'_j}_{s'_j}/\mu^*_{s'_j}} \leq (1+e)^{m(\sigma)}$ In conclusion, we obtain that $\prod_{i=1}^{n}(\rho^{k_i}_{s_i}\rho^{k'_i}_{s'_i}) \leq (1+e)^{m(\sigma)}$. □

6.3 Proof of the sufficiency direction of Theorem 2

We need three lemmas to prove the sufficiency direction. The idea behind the argument is the same as in Echenique and Saito (2015). We know from Lemma 1 that it suffices to find a solution to the relevant system of Afriat inequalities. We take logarithms to linearize the Afriat inequalities in Lemma 1. Then we set up the problem to find a solution to the system of linear inequalities.

The first lemma, Lemma 3, shows that e-PSAROEU is sufficient for e-belief-perturbed OEU rationality under the assumption that the logarithms of the prices are rational numbers. The assumption of rational logarithms comes from our use of a version of the theorem of the alternative (see Lemma 12 in Appendix A.4): when there is no solution to the linearized Afriat inequalities, a rational solution to the dual system of inequalities exists. Then we construct a violation of e-PSAROEU from the given solution to the dual.

The second lemma, Lemma 4, establishes that we can approximate any dataset satisfying e-PSAROEU with a dataset for which the logarithms of prices are rational, and for which e-PSAROEU is satisfied.

The last lemma, Lemma 5, establishes the result by using another version of the theorem of the alternative, stated as Lemma 11.

The rest of the section is devoted to the statement of these lemmas.

Lemma 3. *Given $e \in \mathbf{R}_+$, let a dataset $(x^k, p^k)_{k=1}^{k}$ satisfy e-PSAROEU. Suppose that $\log(p^k_s) \in \mathbf{Q}$ for all $k \in K$ and $s \in S$, $\log(\mu^*_s) \in \mathbf{Q}$ for all $s \in S$, and $\log(1+e) \in \mathbf{Q}$. Then there are numbers v^k_s, λ^k, μ^k_s for $s \in S$ and $k \in K$ satisfying* (10) *and* (11) *in Lemma 1.*

Lemma 4. *Given $e \in \mathbf{R}_+$, let a dataset $(x^k, p^k)_{k=1}^k$ satisfy e-PSAROEU with respect to μ^*. Then for all positive numbers $\bar{\varepsilon}$, there exist a positive real numbers $e' \in [e, e + \bar{\varepsilon}]$, $\mu'_s \in [\mu^*_s - \bar{\varepsilon}, \mu^*_s + \bar{\varepsilon}]$, and $q^k_s \in [p^k_s - \bar{\varepsilon}, p^k_s]$ for all $s \in S$ and $k \in K$ such that $\log q^k_s \in \mathbf{Q}$ for all $s \in S$ and $k \in K$, $\log(\mu'_s) \in \mathbf{Q}$ for all $s \in S$, and $\log(1 + e') \in \mathbf{Q}$, $\mu' \in \Delta_{++}(S)$, and the dataset $(x^k, q^k)_{k=1}^k$ satisfy e'-PSAROEU with respect to μ'.*

Lemma 5. *Given $e \in \mathbf{R}_+$, let a dataset $(x^k, p^k)_{k=1}^k$ satisfy e-PSAROEU with respect to μ. Then there are numbers v^k_s, λ^k, μ^k_s for $s \in S$ and $k \in K$ satisfying* (10) *and* (11) *in Lemma 1.*

6.3.1 Proof of Lemma 3

The proof is similar to the proof of the main result in Echenique and Saito (2015), which corresponds to the case $e = 0$. By log-linearizing the equation in system (10) and the inequality (11) in Lemma 1, we have for all $s \in S$ and $k \in K$, such that

$$\log \mu^k_s + \log v^k_s = \log \lambda^k + \log p^k_s, \tag{12}$$

$$x^k_s > x^{k'}_{s'} \implies \log v^k_s \leq \log v^{k'}_{s'}, \tag{13}$$

and for all $k \in K$ and $s, t \in S$,

$$-\log(1+e) + \log \mu^*_s - \log \mu^*_t \leq \log \mu^k_s - \log \mu^k_t \leq \log(1+e) + \log \mu^*_s - \log \mu^*_t. \tag{14}$$

We are going to write the system of inequalities (12)-(14) in matrix form, following Echenique and Saito (2015) with some modifications.

Let A be a matrix with $K \times S$ rows and $2(K \times S) + K + 1$ columns, defined as follows: We have one row for every pair (k, s), two columns for every pair (k, s), one columns for each k, and one last column. In the row corresponding to (k, s), the matrix has zeroes everywhere with the following exceptions: it has 1's in columns for (k, s); it has a -1 in the column for k; it has $-\log p^k_s$ in the very last column. Matrix A looks as follows:

$$\begin{array}{cc}
 & \begin{array}{cccccccccccccccccc} \cdots & v^k_s & v^k_t & v^l_s & v^l_t & \cdots & \cdots & \mu^k_s & \mu^k_t & \mu^l_s & \mu^l_t & \cdots & \cdots & \lambda^k & \lambda^l & \cdots & p \end{array} \\
\begin{array}{c} \\ (k,s) \\ (k,t) \\ (l,s) \\ (l,t) \\ \\ \end{array} &
\left[\begin{array}{cccccc|cccccc|cccc|c}
 & \vdots & \vdots & \vdots & \vdots & & & \vdots & \vdots & \vdots & \vdots & & & \vdots & \vdots & & \vdots \\
\cdots & 1 & 0 & 0 & 0 & \cdots & \cdots & 1 & 0 & 0 & 0 & \cdots & \cdots & -1 & 0 & \cdots & -\log p^k_s \\
\cdots & 0 & 1 & 0 & 0 & \cdots & \cdots & 0 & 1 & 0 & 0 & \cdots & \cdots & -1 & 0 & \cdots & -\log p^k_t \\
\cdots & 0 & 0 & 1 & 0 & \cdots & \cdots & 0 & 0 & 1 & 0 & \cdots & \cdots & 0 & -1 & \cdots & -\log p^l_s \\
\cdots & 0 & 0 & 0 & 1 & \cdots & \cdots & 0 & 0 & 0 & 1 & \cdots & \cdots & 0 & -1 & \cdots & -\log p^l_t \\
 & \vdots & \vdots & \vdots & \vdots & & & \vdots & \vdots & \vdots & \vdots & & & \vdots & \vdots & & \vdots
\end{array}\right]
\end{array}.$$

Next, we write the system of inequalities (13) and (14) in a matrix form. There is one row in matrix B for each pair (k,s) and (k',s') for which $x^k_s > x^{k'}_{s'}$. In the row corresponding to $x^k_s > x^{k'}_{s'}$, we have zeroes everywhere with the exception of a -1 in the column for (k,s) and a 1 in the column for (k',s'). Matrix B has additional rows, that capture the system of inequalities (14), as follows:

$$
\begin{array}{c}
\begin{array}{cccccc|cccccc|cccc|c}
\cdots & v^k_s & v^k_t & v^l_s & v^l_t & \cdots & \cdots & \mu^k_s & \mu^k_t & \mu^l_s & \mu^l_t & \cdots & \cdots & \lambda^k & \lambda^l & \cdots & p
\end{array} \\
\left[\begin{array}{cccccc|cccccc|cccc|c}
 & \vdots & \vdots & \vdots & \vdots & & & \vdots & \vdots & \vdots & \vdots & & & \vdots & \vdots & & \vdots \\
\cdots & 0 & 0 & 0 & 0 & \cdots & \cdots & 1 & -1 & 0 & 0 & \cdots & \cdots & 0 & 0 & \cdots & \log(1+e) - \log\mu^*_s + \log\mu^*_t \\
\cdots & 0 & 0 & 0 & 0 & \cdots & \cdots & -1 & 1 & 0 & 0 & \cdots & \cdots & 0 & 0 & \cdots & \log(1+e) + \log\mu^*_s - \log\mu^*_t \\
\cdots & 0 & 0 & 0 & 0 & \cdots & \cdots & 0 & 0 & -1 & 1 & \cdots & \cdots & 0 & 0 & \cdots & \log(1+e) + \log\mu^*_s - \log\mu^*_t \\
\cdots & 0 & 0 & 0 & 0 & \cdots & \cdots & 0 & 0 & 1 & -1 & \cdots & \cdots & 0 & 0 & \cdots & \log(1+e) - \log\mu^*_s + \log\mu^*_t \\
 & \vdots & \vdots & \vdots & \vdots & & & \vdots & \vdots & \vdots & \vdots & & & \vdots & \vdots & & \vdots
\end{array}\right]
\end{array}.
$$

Finally, we have a matrix E which has a single row and has zeroes everywhere except for 1 in the last column.

To sum up, there is a solution to the system (12)-(14) if and only if there is a vector $u \in \mathbf{R}^{2(K\times S)+K+1}$ that solves the system of equations and linear inequalities

$$
S1: \begin{cases} A \cdot u = 0, \\ B \cdot u \geq 0, \\ E \cdot u > 0. \end{cases}
$$

The entries of A, B, and E are either 0, 1 or -1, with the exception of the last column of A and B. Under the hypotheses of the lemma we are proving, the last column consists of rational numbers. By Motzkin's theorem, then, there is such a solution u to $S1$ if and only if there is no rational vector (θ,η,π) that solves the system of equations and linear inequalities

$$
S2: \begin{cases} \theta \cdot A + \eta \cdot B + \pi \cdot E = 0, \\ \eta \geq 0, \\ \pi > 0. \end{cases}
$$

In the following, we shall prove that the non-existence of a solution u implies that the dataset must violate e-PSAROEU. Suppose then that there is no solution u and let (θ,η,π) be a rational vector as above, solving system $S2$.

The outline of the rest of the proof is similar to the proof of Echenique and Saito (2015). Since (θ,η,π) are rational vectors, by multiplying a large enough integer, we can make the vectors integers. Then we transform the matrices A and B using θ and η. (i) If $\theta_r > 0$, then creat θ_r copies of the rth row; (ii) omitting row r when $\theta_r = 0$; and (iii) if $\theta_r < 0$, then θ_r copies of the rth row multiplied by -1.

Similarly, we create a new matrix by including the same columns as B and η_r copies of each row (and thus omitting row r when $\eta_r = 0$; recall that $\eta_r \geq 0$ for all r).

By using the transformed matrices and the fact that $\theta \cdot A + \eta \cdot B + \pi \cdot E = 0$ and $\eta \geq 0$, we can prove the following claims:

Claim There exists a sequence $(x^{k_i}_{s_i}, x^{k'_i}_{s'_i})^{n^*}_{i=1} \equiv \sigma$ of pairs that satisfies conditions (i) and (ii) in Definition 6.

Proof. We can construct a sequence $(x^{k_i}_{s_i}, x^{k'_i}_{s'_i})^{n^*}_{i=1}$ in a similar way to the proof of Lemma 11 of Echenique and Saito (2015). By construction, the sequence satisfies condition (i) that $x^{k_i}_{s_i} > x^{k'_i}_{s'_i}$ for all i.

In the following, we show that the sequence satisfies condition (ii) that each k appears as k_i the same number of times it appears as k'_i. Let $n(x^k_s) \equiv \#\{i \mid x^k_s = x^{k_i}_{s_i}\}$ and $n'(x^k_s) \equiv \#\{i \mid x^k_s = x^{k'_i}_{s'_i}\}$. It suffices to show that for each $k \in K$, $\sum_{s\in S}\left[n(x^k_s) - n'(x^k_s)\right] = 0$.

Recall our construction of the matrix B. We have a constraint for each triple (k, s, t) with $s < t$. Denote the weight on the rows capturing $\frac{\mu^k_s/\mu^k_t}{\mu^*_s/\mu^*_t} \leq 1 + e$ by $\eta(k, s, t)$ and $1 + e \leq \frac{\mu^k_s/\mu^k_t}{\mu^*_s/\mu^*_t}$ by $\eta(k, t, s)$.

For each $k \in K$ and $s \in S$, in the column corresponding to μ^k_s in matrix A, remember that we have 1 if we have $x^k_s = x^{k_i}_{s_i}$ for some i and -1 if we have $x^k_s = x^{k'_i}_{s'_i}$ for some i. This is because a row in A must have 1 (-1) in the column corresponding to v^k_s if and only if it has 1 (-1, respectively) in the column corresponding to μ^k_s. By summing over the column corresponding to μ^k_s, we have $n(x^k_s) - n'(x^k_s)$.

Now we consider matrix B. In the column corresponding to μ^k_s, we have 1 in the row multiplied by $\eta(k, t, s)$ and -1 in the row multiplied by $\eta(k, s, t)$. By summing over the column corresponding to μ^k_s, we also have $-\sum_{t\neq s}\eta(k, s, t) + \sum_{t\neq s}\eta(k, t, s)$.

For each $k \in K$ and $s \in S$, the column corresponding to μ^k_s of matrices A and B must sum up to zero; so we have

$$n(x^k_s) - n'(x^k_s) + \sum_{t\neq s}\left[-\eta(k, s, t) + \eta(k, t, s)\right] = 0. \tag{15}$$

Hence for each $k \in K$ for each $k \in K$ $\sum_{s\in S}\left[n(x^k_s) - n'(x^k_s)\right] = 0$. □

Claim $\prod^{n^*}_{i=1}(\rho^{k_i}_{s_i}/\rho^{k'_i}_{s'_i}) > (1 + e)^{m(\sigma^*)}$.

Proof. By (15), So for each $s \in S$

$$\sum_{k\in K}\sum_{s\in S}\sum_{t\neq s}[\eta(k,s,t)-\eta(k,t,s)]\log\mu^*_s=\sum_{k\in K}\sum_{s\in S}\left[n(x^k_s)-n'(x^k_s)\right]\log\mu^*_s=\sum_{i=1}^{n^*}\log\frac{\mu^*_{s_i}}{\mu^*_{s'_i}},$$

where the last equality holds by the definition of n and n'. Moreover, since $d(\sigma^*,k,s)=n(x^k_s)-n'(x^k_s)=\sum_{t\neq s}[\eta(k,s,t)-\eta(k,t,s)]\leq\sum_{t\neq s}\eta(k,s,t)$, we have

$$m(\sigma^*)\equiv\sum_{s\in S}\sum_{k\in K:d(\sigma^*,k,s)>0}d(\sigma^*,k,s)=\sum_{s\in S}\sum_{k\in K}\min\{n(x^k_s)-n'(x^k_s),0\}\leq\sum_{s\in S}\sum_{k\in K}\sum_{t\neq s}\eta(k,s,t).$$

By the equality and the inequality above and by the fact that the last column must sum up to zero and E has one at the last column, we have

$$\begin{aligned}
0&>\sum_{i=1}^{n^*}\log\frac{p^{k'_i}_{s'_i}}{p^{k_i}_{s_i}}+\log(1+e)\sum_{k\in K}\sum_{s\in S}\sum_{t\neq s}\eta(k,s,t)+\sum_{k\in K}\sum_{s\in S}\sum_{t\neq s}(\eta(k,s,t)-\eta(k,t,s))\log\mu^*_s\\
&=\sum_{i=1}^{n^*}\log\frac{p^{k'_i}_{s'_i}}{p^{k_i}_{s_i}}-\sum_{i=1}^{n^*}\log\frac{\mu^*_{s_i}}{\mu^*_{s'_i}}+\log(1+e)\sum_{k\in K}\sum_{s\in S}\sum_{t\neq s}\eta(k,s,t)\\
&=\sum_{i=1}^{n^*}\log\frac{\rho^{k'_i}_{s'_i}}{\rho^{k_i}_{s_i}}+\log(1+e)\sum_{k\in K}\sum_{s\in S}\sum_{t\neq s}\eta(k,s,t)\geq\sum_{i=1}^{n^*}\log\frac{\rho^{k'_i}_{s'_i}}{\rho^{k_i}_{s_i}}+\log(1+e)m(\sigma^*).
\end{aligned}$$

That is, $\sum_{i=1}^{n^*}\log(\rho^{k_i}_{s_i}/\rho^{k'_i}_{s'_i})>m(\sigma^*)\log(1+e)$. This is a contradiction. □

6.3.2 Proof of Lemma 4

Let $\mathcal{X}=\{x^k_s \mid k\in K, s\in S\}$. Consider the set of sequences that satisfy conditions (i) and (ii) in Definition 6:

$$\Sigma=\left\{(x^{k_i}_{s_i},x^{k'_i}_{s'_i})_{i=1}^n\subset\mathcal{X}^2\;\middle|\;\begin{array}{l}(x^{k_i}_{s_i},x^{k'_i}_{s'_i})_{i=1}^n\text{ satisfies conditions (i) and (ii)}\\ \text{in Definition 6 for some } n\end{array}\right\}.$$

For each sequence $\sigma\in\Sigma$, we define a vector $t_\sigma\in\mathbf{N}^{K^2S^2}$. For each pair $(x^{k_i}_{s_i},x^{k'_i}_{s'_i})$, we shall identify the pair with $((k_i,s_i),(k'_i,s'_i))$. Let $t_\sigma((k,s),(k',s'))$ be the number of times that the pair $(x^k_s,x^{k'}_{s'})$ appears in the sequence σ. One can then describe the satisfaction of e-PSAROEU by means of the vectors t_σ. Observe that t depends only on $(x^k)_{k=1}^K$ in the dataset $(x^k,p^k)_{k=1}^K$. It does not depend on prices.

For each $((k,s),(k',s'))$ such that $x^k_s > x^{k'}_{s'}$, define $\delta((k,s),(k',s')) = \log(p^k_s/p^{k'}_{s'})$. And define $\delta((k,s),(k',s')) = 0$ when $x^k_s \le x^{k'}_{s'}$. Then, δ is a K^2S^2-dimensional real-valued vector. If $\sigma = (x^{k_i}_{s_i}, x^{k'_i}_{s'_i})_{i=1}^n$, then

$$\delta \cdot t_\sigma = \sum_{((k,s),(k',s'))\in(KS)^2} \delta((k,s),(k',s'))t_\sigma((k,s),(k',s')) = \log\left(\prod_{i=1}^n \frac{\rho^{k_i}_{s_i}}{\rho^{k'_i}_{s'_i}}\right).$$

So the dataset satisfies e-PSAROEU with respect to μ if and only if $\delta \cdot t_\sigma \le m(\sigma)\log(1+e)$ for all $\sigma \in \Sigma$.

Enumerate the elements in $\mathcal{X}$ in increasing order: $y_1 < y_2 < \cdots < y_N$, and fix an arbitrary $\underline{\xi} \in (0,1)$. We shall construct by induction a sequence $\{(\varepsilon^k_s(n))\}_{n=1}^N$, where $\varepsilon^k_s(n)$ is defined for all (k,s) with $x^k_s = y_n$.

By the denseness of the rational numbers, and the continuity of the exponential function, for each (k,s) such that $x^k_s = y_1$, there exists a positive number $\varepsilon^k_s(1)$ such that $\log(\rho^k_s\varepsilon^k_s(1)) \in \mathbf{Q}$ and $\underline{\xi} < \varepsilon^k_s(1) < 1$. Let $\varepsilon(1) = \min\{\varepsilon^k_s(1) \mid x^k_s = y_1\}$.

In second place, for each (k,s) such that $x^k_s = y_2$, there exists a positive $\varepsilon^k_s(2)$ such that $\log(\rho^k_s\varepsilon^k_s(2)) \in \mathbf{Q}$ and $\underline{\xi} < \varepsilon^k_s(2) < \varepsilon(1)$. Let $\varepsilon(2) = \min\{\varepsilon^k_s(2) \mid x^k_s = y_2\}$.

In third place, and reasoning by induction, suppose that $\varepsilon(n)$ has been defined and that $\underline{\xi} < \varepsilon(n)$. For each (k,s) such that $x^k_s = y_{n+1}$, let $\varepsilon^k_s(n+1) > 0$ be such that $\log(\rho^k_s\varepsilon^k_s(n+1)) \in \mathbf{Q}$, and $\underline{\xi} < \varepsilon^k_s(n+1) < \varepsilon(n)$. Let $\varepsilon(n+1) = \min\{\varepsilon^k_s(n+1) \mid x^k_s = y_n\}$.

This defines the sequence $(\varepsilon^k_s(n))$ by induction. Note that $\varepsilon^k_s(n+1)/\varepsilon(n) < 1$ for all n. Let $\bar{\xi} < 1$ be such that $\varepsilon^k_s(n+1)/\varepsilon(n) < \bar{\xi}$.

For each $k \in K$ and $s \in S$, let $\hat{\rho}^k_s = \rho^k_s\varepsilon^k_s(n)$, where n is such that $x^k_s = y_n$. Choose $\mu' \in \Delta_{++}(S)$ such that for all $s \in S$ $\log\mu'_s \in \mathbf{Q}$ and $\mu'_s \in [\bar{\xi}\mu_s, \mu_s/\bar{\xi}]$ for all $s \in S$. Such μ' exists by the denseness of the rational numbers. Now for each $k \in K$ and $s \in S$, define

$$q^k_s = \frac{\hat{\rho}^k_s}{\mu'_s}. \tag{16}$$

Then, $\log q^k_s = \log\hat{\rho}^k_s - \log\mu'_s \in \mathbf{Q}$.

We claim that the dataset $(x^k, q^k)_{k=1}^K$ satisfies e'-PSAROEU with respect to μ'. Let δ^* be defined from $(q^k)_{k=1}^K$ in the same manner as δ was defined from $(\rho^k)_{k=1}^K$.

For each pair $((k,s),(k',s'))$ with $x^k_s > x^{k'}_{s'}$, if n and m are such that $x^k_s = y_n$ and $x^{k'}_{s'} = y_m$, then $n > m$. By definition of ε,

$$\frac{\varepsilon^k_s(n)}{\varepsilon^{k'}_{s'}(m)} < \frac{\varepsilon^k_s(n)}{\varepsilon(m)} < \bar{\xi} < 1.$$

Hence,

$$\delta^*((k,s),(k',s')) = \log \frac{\rho_s^k \varepsilon_s^k(n)}{\rho_{s'}^{k'} \varepsilon_{s'}^{k'}(m)} < \log \frac{\rho_s^k}{\rho_{s'}^{k'}} + \log \bar{\xi} < \log \frac{\rho_s^k}{\rho_{s'}^{k'}} = \delta((k,s),(k',s')).$$

Now, we choose e' such that $e' \geq e$ and $\log(1+e') \in \mathbf{Q}$.

Thus, for all $\sigma \in \Sigma$, $\delta^* \cdot t_\sigma \leq \delta \cdot t_\sigma \leq m(\sigma)\log(1+e) \leq m(\sigma)\log(1+e')$ as $t. \geq 0$ and the dataset $(x^k, p^k)_{k=1}^K$ satisfies e-PSAROEU with respect to μ.

Thus the dataset $(x^k, q^k)_{k=1}^K$ satisfies e'-PSAROEU with respect to μ'. Finally, note that $\underline{\xi} < \varepsilon_s^k(n) < 1$ for all n and each $k \in K, s \in S$. So that by choosing $\underline{\xi}$ close enough to 1, we can take $\hat{\rho}$ to be as close to ρ as desired. By the definition, we also can take μ' to be as close to μ as desired. Consequently, by (16), we can take $(q^k)_{k=1}^K$ to be as close to $(p^k)_{k=1}^K$ as desired. We also can take e' to be as close to e as desired.

6.3.3 Proof of Lemma 5

We use the following notational convention: For a matrix D with $2(K \times S) + K + 1$ columns, write D_1 for the submatrix of D corresponding to the first $K \times S$ columns; let D_2 be the submatrix corresponding to the following $K \times S$ columns; D_3 correspond to the next K columns; and D_4 to the last column. Thus, $D = [D_1|D_2|D_3|D_4]$.

Consider the system comprised by (12), (13), and (14) in the proof of Lemma 3. Let A, B, and E be constructed from the dataset as in the proof of Lemma 3. The difference with respect to Lemma 3 is that now the entries of A_4 and B_4 may not be rational. Note that the entries of E, B, and A_i, $i = 1, 2, 3$ are rational.

Suppose, towards a contradiction, that there is no solution to the system comprised by (12), (13), and (14). Then, by the argument in the proof of Lemma 3 there is no solution to system $S1$. Lemma 11 (in Appendix A.4) with $\mathbf{F} = \mathbf{R}$ implies that there is a real vector (θ, η, π) such that $\theta \cdot A + \eta \cdot B + \pi \cdot E = 0$ and $\eta \geq 0, \pi > 0$. Recall that $E_4 = 1$, so we obtain that $\theta \cdot A_4 + \eta \cdot B_4 + \pi = 0$.

Consider $(q^k)_{k=1}^K$, μ', and e' be such that the dataset $(x^k, q^k)_{k=1}^K$ satisfies e'-PSAROEU with respect to μ', and $\log q_s^k \in \mathbf{Q}$ for all k and s, $\log \mu'_s \in \mathbf{Q}$ for all $s \in S$, and $\log(1+e') \in \mathbf{Q}$. (Such $(q^k)_{k=1}^K$, μ', and e' exist by Lemma 4.) Construct matrices A', B', and E' from this dataset in the same way as A, B, and E is constructed in the proof of Lemma 3. Note that only the prices, the objective probabilities, and the bounds are different. So $E' = E$ and $A'_i = A_i$ and $B'_i = B_i$ for $i = 1, 2, 3$. Only A'_4 and B'_4 may be different from A_4 and B_4, respectively.

By Lemma 4, we can choose q^k, μ', and e' such that $|(\theta \cdot A'_4 + \eta \cdot B'_4) - (\theta \cdot A_4 + \eta \cdot B_4)| < \pi/2$. We have shown that $\theta \cdot A_4 + \eta \cdot B_4 = -\pi$, so the choice of q^k, μ', and e' guarantees that $\theta \cdot A'_4 + \eta \cdot B'_4 < 0$. Let $\pi' = -\theta \cdot A'_4 - \eta \cdot B'_4 > 0$.

Note that $\theta \cdot A'_i + \eta \cdot B'_i + \pi' E_i = 0$ for $i = 1, 2, 3$, as (θ, η, π) solves system $S2$ for matrices A, B and E, and $A'_i = A_i$, $B'_i = B_i$ and $E_i = 0$ for $i = 1, 2, 3$. Finally, $\theta \cdot A'_4 + \eta \cdot B'_4 + \pi' E_4 = \theta \cdot A'_4 + \eta \cdot B'_4 + \pi' = 0$. We also have that $\eta \geq 0$ and $\pi' > 0$. Therefore θ, η, and π' constitute a solution to $S2$ for matrices A', B', and E'.

Lemma 11 then implies that there is no solution to system $S1$ for matrices A', B', and E'. So there is no solution to the system comprised by (12), (13), and (14) in the proof of Lemma 3. However, this contradicts Lemma 3 because the dataset (x^k, q^k) satisfies e'-PSAROEU with μ', $\log(1 + e') \in \mathbf{Q}$, $\log \mu'_s \in \mathbf{Q}$ for all $s \in S$, and $\log q^k_s \in \mathbf{Q}$ for all $k \in K$ and $s \in S$.

References

Afriat, Sidney N., "The Construction of Utility Functions from Expenditure Data," *International Economic Review*, 1967, *8* (1), 67–77.

— , "Efficiency Estimation of Production Functions," *International Economic Review*, 1972, *13* (3), 568–598.

Ahn, David S., Syngjoo Choi, Douglas Gale, and Shachar Kariv, "Estimating Ambiguity Aversion in a Portfolio Choice Experiment," *Quantitative Economics*, 2014, *5* (2), 195–223.

Allen, Roy and John Rehbeck, "Assessing Misspecification and Aggregation for Structured Preferences," 2018. Unpublished manuscript.

Apesteguia, Jose and Miguel A. Ballester, "A Measure of Rationality and Welfare," *Journal of Political Economy*, 2015, *123* (6), 1278–1310.

Carvalho, Leandro and Dan Silverman, "Complexity and Sophistication," 2017. Unpublished manuscript.

— , Stephan Meier, and Stephanie W. Wang, "Poverty and Economic Decision Making: Evidence from Changes in Financial Resources at Payday," *American Economic Review*, 2016, *106* (2), 260–284.

Chambers, Christopher P. and Federico Echenique, *Revealed Preference Theory*, Cambridge: Cambridge University Press, 2016.

— , Ce Liu, and Seung-Keun Martinez, "A Test for Risk-Averse Expected Utility," *Journal of Economic Theory*, 2016, *163*, 775–785.

Choi, Syngjoo, Raymond Fisman, Douglas Gale, and Shacher Kariv, "Consistency and Heterogeneity of Individual Behavior under Uncertainty," *American Economic Review*, 2007, *97* (5), 1921–1938.

— , Shachar Kariv, Wieland Müller, and Dan Silverman, "Who Is (More) Rational?," *American Economic Review*, 2014, *104* (6), 1518–1550.

de Clippel, Geoffroy and Kareen Rozen, "Relaxed Optimization: ε-Rationalizability and the FOC-Departure Index in Consumer Theory," 2019. Unpublished manuscript.

Dean, Mark and Daniel Martin, "Measuring Rationality with the Minimum Cost of Revealed Preference Violations," *Review of Economics and Statistics*, 2016, *98* (3), 524–534.

Dziewulski, Pawel, "Eliciting the Just-Noticeable Difference," 2016. Unpublished manuscript.

_ , "Just-Noticeable Difference as a Behavioural Foundation of the Critical Cost-Efficiency Index," 2018. Unpublished manuscript.

Echenique, Federico and Kota Saito, "Savage in the Market," *Econometrica*, 2015, *83* (4), 1467–1495.

_ , Sangmok Lee, and Matthew Shum, "The Money Pump as a Measure of Revealed Preference Violations," *Journal of Political Economy*, 2011, *119* (6), 1201–1223.

_ , Taisuke Imai, and Kota Saito, "Testable Implications of Models of Intertemporal Choice: Exponential Discounting and Its Generalizations," 2016. Caltech HSS Working Paper 1388.

Frederick, Shane, "Cognitive Reflection and Decision Making," *Journal of Economic Perspectives*, 2005, *19* (4), 25–42.

Friedman, Daniel, Sameh Habib, Duncan James, and Sean Crockett, "Varieties of Risk Elicitation," 2018. Unpublished manuscript.

Green, Richard C. and Sanjay Srivastava, "Expected Utility Maximization and Demand Behavior," *Journal of Economic Theory*, 1986, *38* (2), 313–323.

Griliches, Zvi, "Economic Data Issues," in Zvi Griliches and Michael D. Intriligator, eds., *Handbook of Econometrics*, Vol. 3, Elsevier, 1986, pp. 1465–1514.

Halevy, Yoram, Dotan Persitz, and Lanny Zrill, "Parametric Recoverability of Preferences," *Journal of Political Economy*, 2018, *126* (4), 1558–1593.

Kubler, Felix, Larry Selden, and Xiao Wei, "Asset Demand Based Tests of Expected Utility Maximization," *American Economic Review*, 2014, *104* (11), 3459–3480.

Loomes, Graham, "Evidence of a New Violation of the Independence Axiom," *Journal of Risk and Uncertainty*, 1991, *4* (1), 91–108.

McFadden, Daniel, "Conditional Logit Analysis of Qualitative Choice Behavior," in Paul Zarembka, ed., *Frontiers in Econometrics*, New York: Academic Press, 1974, pp. 105–142.

Polisson, Matthew, John K.-H. Quah, and Ludovic Renou, "Revealed Preferences over Risk and Uncertainty," 2017. Unpublished manuscript.

Samuelson, Paul A., "A Note on the Pure theory of Consumer's Behaviour," *Economica*, 1938, *5* (17), 61–71.

Varian, Hal R., "The Nonparametric Approach to Demand Analysis," *Econometrica*, 1982, *50* (4), 945–973.

_ , "Goodness-of-Fit in Optimizing Models," *Journal of Econometrics*, 1990, *46* (1), 125–140.