RELAXED OPTIMIZATION:
ε-RATIONALIZABILITY AND THE FOC-DEPARTURE INDEX IN CONSUMER THEORY*

Geoffroy de Clippel† and Kareen Rozen‡

This Version: July 2019

Abstract

We propose relaxing the first-order conditions in optimization to approximate rational consumer choice. Departures from the FOCs are assessed using an axiomatically founded measure that is also interpretable in terms of a money-pump multiplier. The framework encompasses measurement errors, information unobservable to the modeler, and consumer misperception. We develop testable implications for demand data, including for subclasses of regular utility functions, and develop the FOC-Departure Index (FDI), which is applicable in all contexts where the first-order approach is meaningful. We extend these ideas to convex budget sets. Our analysis extends to non-convex preferences under a narrower interpretation of price misperception.

*First version July 2018 (under the title 'Consumer Theory with Misperceived Tastes'). We thank Andrew Postlewaite and seminar participants for helpful comments. We are grateful to Tommaso Coen and Giacomo Rubbini for outstanding research assistance.

†Department of Economics, Brown University. Email address: declippel@brown.edu.

‡Department of Economics, Brown University. Email address: kareen_rozen@brown.edu.

1. Introduction

A central idea behind rationality is the ability to optimize. In a typical setting, this entails choosing an option that balances the marginal benefits and marginal costs along all dimensions. Consistency with rationality leaves no room for departure from these first-order conditions. One could imagine a more nuanced approach, yielding an approximation of rationality that is gradated by the extent of departure permitted. Such an approach naturally begs the question, how should departure from first-order conditions be measured?

In this paper, we propose such an approach in the classic setting of consumer theory, using an axiomatically-based notion of departure from first-order conditions. For the purposes of this introduction only, we present the key idea under the simplifying assumption that utility functions are everywhere differentiable. Given an $\varepsilon \in [0,1]$, we say demand data is ε-rationalizable if there is a utility function u, satisfying some regularity properties[1] (ensuring, in particular, that the second-order condition applies), such that if we observe the bundle x being chosen at a price vector p, then

$$1 - \varepsilon \leq \frac{MRS^u_{\ell\ell'}(x)}{p_\ell / p_{\ell'}} \leq \frac{1}{1-\varepsilon}, \tag{1}$$

for all $\ell \neq \ell'$. The classic first-order condition, which requires marginal rates of substitution (MRS) and opportunity costs to match, corresponds to ε-rationalizability for $\varepsilon = 0$. The larger ε is, the more the MRS may depart from opportunity costs, and the more permissive ε-Rationality becomes. The particular way (1) measures departure from the first-order condition may initially seem ad hoc. To the contrary, we show that for any pair of goods it is essentially the *unique* method of measuring departures between objects such as price vectors and utility gradients, given the economic properties such objects must respect. For instance, notice that (1) is invariant to the units in which goods are measured and priced. We provide further interpretations of our measure of departure as a money-pump multiplier and as a measure of parameter misspecification for additively-separable utility functions.

There may be multiple reasons demand data is inconsistent with perfect optimization. There may be measurement errors when collecting price data, such as unknown promotions, sales, coupons, overhead expenses, or tax implications. Even if prices are accurately recorded, there may be factors that rationally affect choice, but which are

[1] Specifically, we will require continuity, strict monotonicity and strict quasi-concavity.

not included in the dataset. For instance, utility tradeoffs may vary a bit based on expiration dates or ripeness of produce, which are observed by the consumer but not the modeler. Bounded rationality may play a role as well. The consumer could have an imperfect understanding of the prices, or misperceive her utility tradeoffs.[2] This misperception could be unintended or due to active framing by third parties (e.g., a persuasive salesperson, packaging, store displays). Our framework accommodates such data, and encompasses all these interpretations. More generally, the measure is of interest for quantifying how much choices depart from being rationalizable, independently of the source of the departure.

We develop the testable implications of ε-rationalizability, showing it is equivalent to a relaxation of SARP. Namely, there must be an acyclic preference over the finite set of demanded bundles, satisfying certain restrictions (which are weaker than Samuelson's revealed preference) on which demanded bundles are worse than others. We then develop interesting properties of the measure of choice inconsistency to which ε-rationalizability naturally lends itself. Namely, the FOC-Departure Index (FDI) of demand data is defined as the infimum over all ε for which the data is ε-rationalizable. We provide a simple way to compute the FDI in general, by solving linear programming problems. Better yet, in the case of two-dimensional bundles (the setting of most experiments), we show the FDI is given by the maximum FDI from each pair of observations, which is itself computed using an easy, closed-form formula.

The FOC-Departure Index measures the smallest departure from the first-order conditions with which the data is consistent. By contrast, the classic Critical Cost Efficiency Index of Afriat (1973) considers global monetary effects, measuring the percentage of income that can be retained while eliminating revealed preference cycles. Despite these differences, we show that there is a surprising relationship between the two measures: the FDI is bounded below by the percent of income lost according to the CCEI (that is, $1 - \text{CCEI} \leq \text{FDI}$ for all datasets). Phrased differently, small departures from the first-order conditions imply only small budgetary adjustments are needed to eliminate revealed preference cycles, but not vice-versa. This would suggest that our measure is more demanding than that of Afriat, but the story is subtler once power is considered: one should take into account whether violations of rationality are

[2]The topic of perception has garnered recent interest in the economic literature; see, among other works, Gabaix (2014) on perceiving prices optimally, Woodford (2012) and Steiner and Stewart (2016) on perceiving risky prospects, and Esponda (2016) on an equilibrium framework for agents who misperceive their environment.

likely for the budget sets observed. Using Bronars (1987)'s well-known approach,[3] for instance, we show that there exist datasets where the FDI suggests greater rationality than does the CCEI, as well as datasets where the opposite is true.

In many contexts, one may want to understand ε-rationalizability and the FOC-Departure Index while restricting to a special class of regular utility functions (e.g., also requiring quasi-linearity, additive separability, homotheticity, or exponential discounting). We extend these notions to *any* subclass $\mathcal{U}$ of regular utility functions, developing important connections with Afriat's CCEI and classic tests of rationalizability based on Afriat-like inequalities. First, we show that the inequality $1 - \text{CCEI} \leq \text{FDI}$ that we uncovered continues to hold when restricting the computations of the CCEI and FDI to any subclass $\mathcal{U}$. This provides a bound on the CCEI which may be of independent interest, since it is not always obvious how to compute it with respect to special classes. Second, we provide a single umbrella result that shows, for any subclass $\mathcal{U}$ of regular utility functions, how to test ε-rationalizability and compute the FDI by repurposing any existing test of rationalizability with respect to $\mathcal{U}$ that is based on Afriat-like inequalities. This is quite convenient for applications.

We illustrate the applicability of our approach using the portfolio-choice dataset of Choi, Kariv, Müller and Silverman (2014), which consists of 1,182 adults recruited from the CentERpanel sample. In their experiment, subjects make decisions for 25 randomly-drawn, two-dimensional budget sets, where a bundle (x_1, x_2) describes the monetary payment in each of two equally-likely states. We begin by examining this data using the FDI, rather than the CCEI as in Choi et al. (2014). Though correlated, the measures are notably different. Not only is $\text{FDI} > 1 - \text{CCEI}$ for over 75% of subjects, but the two measures also suggest, in 15% of all subject pairings, opposite rankings of who is more rational. Nonetheless, an exercise in the style of Bronars (1987) confirms Choi et al's observation that there is a significant amount of rationality to be found. We then take advantage of the FDI's flexibility to study departures from subclasses of regular utility functions: besides expected utility, we also examine the wider class of state independent utility. Indeed, as the two states are equally likely and have no underlying meaning, the objects of choice are effectively lotteries; and all preferences over lotteries in this setting would deem (x_1, x_2) as good as (x_2, x_1), as both induce an equal chance of winning \$$x_1$ or \$$x_2$.[4] In a contemporaneous and

[3]The idea is to compare the distribution of the index under the true data to the distribution arising if choices were drawn uniformly from budget frontiers.

[4]With equally-likely states and monotonicity with respect to money, the symmetry that follows

independent paper, Echenique, Imai and Saito (2019) define a measure of departure from risk-averse expected utility that essentially coincides with the restriction of the FDI to that subclass of preferences. They apply this measure to several previously-collected experimental datasets, including Choi et al., to examine the prevalence of expected utility. Our ability to vary the reference class of preferences under the same measure of departure allows us to add further insight. Perhaps surprisingly, we find that the empirical departure from expected utility is mostly attributable to departure from the natural property of state independence, even though expected utility is theoretically much more demanding. The underlying issue is that while the dataset is well-suited for testing basic rationality, a richer dataset (e.g., varying the probability of states) is needed to better disentangle expected utility from the plain maximization of a preference over lotteries.

A benefit of our approach is that it applies whenever the first-order approach is meaningful. We take advantage of this flexibility to extend ε-rationalizability and the FDI to generalized demand data with non-linear pricing. Notice that applying Afriat's index requires a context in which 'shrinking the choice set' is meaningful. This is unclear with non-linear budget sets, but poses no issue for ε-rationalizability and the FDI, which leave choice sets unchanged. The feasible set must be convex for the extended notion to be interpretable in terms of departure from global optimality. We show that one can reduce all tests and computations for demand data with non-linear pricing to tests and computations over a fictitious dataset with *linear* prices, given by the marginal prices at observed demands (or a generalization of marginal prices at non-differentiable points of the pricing schedule). We illustrate the approach with a simple example of progressive taxation.

The use of first-order conditions is central to many parts of economics, and naturally, the second-order condition is assumed for it to be meaningful, that is, ensure a global optimum. Without assuming quasi-concavity of the utility function, any demand data would be ε-rationalizable (for all ε): while the inequalities in (1) do restrict utility gradients given prices, one can always draw indifference curves which match these local restrictions with only finitely many observations. It may come as a surprise, therefore, that demand data is ε-rationalizable if and only if there is a strictly monotone and continuous utility function u (not required to be quasi-concave!) such that for each observation (p, x), the bundle x is u-maximal in the budget set associated

from state independence is also tantamount to respecting first-order stochastic dominance.

to the price vector p^c, with

$$1-\varepsilon < \frac{p^c_\ell/p^c_{\ell'}}{p_\ell/p_{\ell'}} < \frac{1}{1-\varepsilon},$$

satisfied for all ℓ, ℓ'. Here p represents the price vector collected by the modeler, while p^c is the price vector used by the consumer.[5] These prices need not coincide: the consumer may use the correct prices while the modeler's record is faulty; the consumer may misunderstand the true prices while the modeler's record is accurate; or some combination of these scenarios may hold. In all cases, if the presumption is that the consumer optimizes correctly, then ϵ-rationalizability admits a natural interpretation that *does not rely on quasi-concavity.* This is in the spirit of Afriat (1967)'s result that quasi-concavity is not testable under the rational choice model.

We conclude the paper by further developing the idea of a Price-Misperception Index (PMI), including for special classes of utility functions. Though narrower in interpretation, such an extension of the FDI to broader classes of utility functions has the benefit of not relying on quasi-concavity. The PMI coincides with the FDI with respect to classes of regular utility functions. We show that to find the PMI more generally, we can similarly repurpose standard rationalizability tests (e.g., for expected utility without requiring risk aversion, or other interesting classes). We revisit our analysis of Choi et al. (2014) using the PMI, which reinforces our earlier conclusions even when allowing for more general risk preferences.

The paper is organized as follows. We formalize the framework in Section 2. There, we consider the axiomatic basis of our notion of departure from first-order conditions; and provide further rationales in terms of price mismatch, money pumps in the case of price misperception, and parameter misspecification in the common setting of additively separable utility. In Section 3, we characterize the testable implications for the class of regular utility functions. In Section 4, we develop the FDI and its properties. In Section 5, we further develop these ideas for arbitrary subclasses of regular utility functions, and illustrate our methodology with an application to experimental data where bundles correspond to lotteries. We extend beyond linear budget sets in Section 6, and beyond convex preferences (i.e., quasi-concave utility functions) in Section 7.

[5] Our notion of departure thus applies here to the price vectors (the axiomatic analysis covers any combination of objects such as gradients or price vectors).

2. Consumer Data and ε-Rationalizability

We observe a consumer selecting a consumption bundle at various price vectors. The *demand data* $\mathcal{D}$ comprises a finite collection of pairs (p, x), where $p \in \mathbb{R}^L_{++}$ is a *price vector* and $x \in \mathbb{R}^L_+$ is the *consumption bundle* demanded at p.

As usual, preference orderings will be assumed to be continuous and strictly monotone. We further assume strict convexity (we refer the reader to Footnote 19 on relaxing strictness, and Section 7 on relaxing convexity). Such preferences are representable by a *regular* utility function, namely one that is continuous, strictly monotone, and strictly quasi-concave. The rational benchmark posits that the consumer selects bundles through utility maximization over budget sets. Should the utility function be differentiable at an interior choice, opportunity costs must equal marginal rates of substitution. Without requiring full differentiability or interior choices, optimality means each price vector must separate the upper-contour set of the choice from all other affordable bundles.

Formally, given a utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$, define $\partial u(x)$ to be the set of strictly positive vectors defining the supporting hyperplanes of the upper-contour set at x:[6]

$$\partial u(x) = \{g \in \mathbb{R}^L_{++} \mid \forall y : u(y) \geq u(x) \Rightarrow g \cdot y \geq g \cdot x\}.$$

This set is economically important: it captures the local shape of the consumer's indifference curves, by characterizing quantity tradeoffs that leave the consumer indifferent. Elements of $\partial u(x)$ are called *quasi-gradients*. At any point x where u is differentiable, $\partial u(x)$ contains a unique vector – the usual gradient $\nabla u(x) = (\frac{\partial u(x)}{\partial x_1}, \ldots, \frac{\partial u(x)}{\partial x_L})$ – up to positive rescaling.

Definition 1 (ε-Rationalizability) *For $\varepsilon \in [0, 1]$, the demand data $\mathcal{D}$ is* ε-rationalizable *if there exists a regular utility function u and, for each $(p, x) \in \mathcal{D}$, a vector $g \in \partial u(x)$ such that*

$$1 - \varepsilon \leq \frac{g_\ell / g_{\ell'}}{p_\ell / p_{\ell'}} \leq \frac{1}{1 - \varepsilon}, \tag{2}$$

for each $\ell \neq \ell'$. Then u is said to ε-rationalize *the demand data.*

[6] At the cost of heavier notation later on, one could allow g to belong to $\mathbb{R}^I_+ \setminus \{0\}$ in the definition of $\partial u(x)$ (a change that matters only if some component of x is zero, by strict monotonicity). This makes no difference in any of our results because the added vectors are always as far as it gets, according to our measure of departure below, to strictly positive price vectors.

For further intuition, the ratio $g_\ell/g_{\ell'}$ is uniquely defined for $x \gg 0$ when u is differentiable, and corresponds to the *marginal rate of substitution* (MRS) associated to ℓ and ℓ':

$$MRS^u_{\ell\ell'}(x) = \frac{\partial u(x)/\partial x_\ell}{\partial u(x)/\partial x_{\ell'}}.$$

The numerator in (2) simply captures tradeoffs in taste, while accommodating the possibility of multiple implicit utility tradeoffs in case of non-differentiability. By contrast, the price ratio $p_\ell/p_{\ell'}$ in the denominator represents an opportunity cost. Notice then that (2) simply boils down to the FOC that is characteristic of rationality when $\varepsilon = 0$. There are several reasons to be interested in ε-rationalizability as a relaxation of rationality, as explained in the following subsections.

2.1 How Far Apart Must Utility Gradients and Price Vectors Be?

Suppose that demand data appears inconsistent with rationality. Can we quantify how close a miss it was? Rationality holds if and only if price vectors are quasi-gradients of the utility function (or price vectors and utility gradients are collinear in case of differentiability). Our suggestion then is simply to measure how far apart utility gradients or quasi-gradients must be from price vectors, while considering all possible utility functions. But how does one measure the degree of discrepancy between any two such vectors in the first place?

Consider the simpler case of two goods first. Some obvious candidates, such as the angular distance or the Euclidean distance, pose an issue when the vectors are understood to be prices or marginal utilities/quasi-gradients. Indeed, given their economic interpretation, such variables are not uniquely defined. For instance, an increasing transformation of a utility function offers another representation of the same preference, with rescaled marginal utilities; and similarly, rescaling prices leaves budget sets unchanged. The Euclidean distance, however, is sensitive to such rescaling. Moreover, the consumer's problem is fundamentally unaffected when modifying how any given good's quantity is measured (e.g., using ounces or grams, gallons or quarts, etc.), provided that prices are adjusted accordingly.[7] Both angular and Euclidean distance, however, are sensitive to measurement choice.

[7] For instance, buying one gallon of milk at \$4, is the same as buying 4 quarts at a quarter of the price (\$1 each). Also, the marginal utility from buying an extra η quarts, is a quarter of the marginal utility from buying an extra η gallons (where $\eta > 0$ is small).

Let $\succeq$ be a weak ordering over vector pairs $(x, y) \in \mathbb{R}^2_{++}$: $(x, y) \succeq (x', y')$ means that "x is farther from y than x' is from y'." For concreteness, it may help to think of x, x' as utility gradients and y, y' as price vectors. The following axioms capture the invariance properties discussed above.

Unit Invariance $(x, y) \succeq (x', y')$ *if and only if* $(\alpha x, \beta y) \succeq (\alpha' x', \beta' y')$, *for all positive scalars* $\alpha, \beta, \alpha', \beta'$.

Measurement Invariance $(x, y) \succeq (x', y')$ *if and only if* $((\alpha x_1, x_2), (\alpha y_1, y_2)) \succeq (\alpha' x'_1, x'_2), (\alpha' y'_1, y'_2))$, *for all positive scalars* α, α' *(and similarly for good* 2*).*

The first axiom reflects the fact that a price vector or a utility (quasi-)gradient is effectively determined only up to a positive linear transformation. It permits different transformations for the different vectors x, x', y and y', but requires *all dimensions* of any given vector to be scaled by the same factor. The second axiom considers a different type of transformation, whereby the *same dimension* is scaled by the same factor in the pair of vectors compared. It captures invariance to the way in which we measure a good, and thus also how we state its price.

These two invariance properties go a long way in determining the structure of $\succeq$. Indeed, adding only the following three regularity properties uniquely pins it down.

Representability $\succeq$ *is complete, transitive and continuous.*

Symmetry $(x, y) \sim (y, x)$

Monotonicity $((\hat{\alpha}, 1), (1, 1)) \succ ((\alpha, 1), (1, 1))$ *for all* $1 \leq \alpha < \hat{\alpha}$.

The first axiom ensures existence of a numerical representation. The second means that x is equally far from y as y is from x. The third simply requires that an increase in $\alpha \geq 1$ brings the vector $(\alpha, 1)$ further from $(1, 1)$.[8]

Proposition 1 *There is a unique ordering* $\succeq^*$ *satisfying the five axioms:* $(x', y') \succeq^* (x, y)$ *if and only if* $\delta(x', y') \geq \delta(x, y)$, *where*

$$\delta(a, b) = \max\{\frac{a_1/a_2}{b_1/b_2}, \frac{b_1/b_2}{a_1/a_2}\},$$

for all vectors a, b *in* $\mathbb{R}^2_{++}$.

The representation in Proposition 1 compares how far apart two vectors are, in

[8]Naturally, one would also desire the opposite relationship in the case $\alpha \in (0, 1)$. Imposing this property would be redundant, however: the ordering we uncover in Proposition 1 satisfies it.

comparison to how far apart another two reference vectors are. The choice of the reference vectors simply parametrizes a bound. To fix ideas, one can measure how far x is from y, relative to how far $(1,1)$ is from $(1-\varepsilon,1)$ for $\varepsilon \in [0,1)$. In that case:

$$\delta((1-\varepsilon,1),(1,1)) = \max\{1-\varepsilon, \frac{1}{1-\varepsilon}\} = \frac{1}{1-\varepsilon}.$$

Hence $((1-\varepsilon,1),(1,1)) \succeq (x,y)$ if and only if

$$\delta(x,y) = \max\{\frac{x_1/x_2}{y_1/y_2}, \frac{y_1/y_2}{x_1/x_2}\} \leq \frac{1}{1-\varepsilon},$$

which is equivalent to how Definition 1 of ε-rationalizability measures how far apart are x and y, where one is a utility quasi-gradient and the other is a price vector:

$$1-\varepsilon \leq \frac{x_1/x_2}{y_1/y_2} \leq \frac{1}{1-\varepsilon}.$$

To extend beyond two goods, one can for instance consider the projections of these vectors to each pairs of goods, and take the largest distance:

$$\delta(x,y) = \max_{\ell,\ell'} \delta((x_\ell, x_{\ell'}),(y_\ell, y_{\ell'})).$$

Our axiomatic characterization thus suggests ε-rationalizability is a natural way to quantify the extent to which the consumer deviates from the optimal tradeoffs in her first-order condition, while remaining agnostic about the source of this deviation. Note also that the axioms equally apply whether we use the measure to evaluate how far apart are two price vectors, how far apart are two utility quasi-gradients, or how far apart is a price vector from a utility quasi-gradient.

2.2 Misperceived Prices

As we have discussed, the prices p on record and the prices p^c the consumer takes into account need not coincide. The consumer may use the correct prices while the modeler's record is faulty; the consumer may misunderstand the true prices while the modeler's record is accurate; or some combination of these scenarios may hold. As the next proposition shows, ε-rationalizability is equivalent to the consumer picking the optimal bundles from what she understands the budget set to be, under the requirement that her prices are not too far from the modeler's, using the same δ-criterion uncovered axiomatically in the previous subsection.

Proposition 2 *$\mathcal{D}$ is ε-rationalizable if and only if there exists a strictly monotone and continuous utility function $v : \mathbb{R}^L_+ \to \mathbb{R}$ such that one can associate to all $(p, x) \in \mathcal{D}$ a price vector $p^c \in \mathbb{R}^L_{++}$ used by the consumer with*

$$x \in \operatorname*{argmax}_{\{p^c \cdot y \leq p^c \cdot x\}} v(y), \ \textit{and} \tag{3a}$$

$$1 - \varepsilon < \frac{p_\ell / p_{\ell'}}{p^c_\ell / p^c_{\ell'}} < \frac{1}{1-\varepsilon}, \tag{3b}$$

for all $\ell \neq \ell'$.

Necessity is rather straightforward. Indeed, by ε-rationalizability, there exists a regular utility function u for which ∂u admits a selection that is never too far apart from price vectors, for observations in $\mathcal{D}$. The result then follows almost immediately by taking these vectors as the consumer's price vectors. Note though that, following (2), we get (3b) only with weak inequalities. We prove in Lemma 3 of the Appendix that, perhaps surprisingly, requiring weak or strict inequalities in Definition 1 makes no difference. With this additional observation, the proof of necessity is complete. Sufficiency would be rather straightforward too if v were required to be regular, but notice that Proposition 2 does *not* impose any quasi-concavity on v. This makes the result quite interesting, but also harder to prove.[9] So, while the definition of ε-rationalizability relies on regular utility functions, the foundation for it in terms of price mismatch does not.

Under an interpretation of price misperception by the consumer, Proposition 2 captures a notion of bounded rationality that is related to Gabaix (2014), but somewhat different in spirit. The consumer picks the optimal bundle in her perceived budget set in both cases. However, Gabaix proposes a theory of how perceived prices arise as a function of an exogenous default price vector and the actual price vector, while the theory arising in Proposition 2 only requires the perceived price vector to be in the vicinity of the true price vector.

[9] Afriat (1967)'s result that quasi-concavity has no empirical content with linear budget sets is of no help, as it does not extend to strict quasi-concavity. Instead, we prove in the Appendix (see Lemma 2) that any dataset satisfying GARP admits a nearby dataset, varying only prices as little as desired, that satisfies SARP.

Beyond the axiomatic foundation in Section 2.1, Equation (3b) admits an additional interpretation when the mismatch between prices arises from the consumer's misperception. The discrepency between p and p^c means the consumer can be subject to a money pump. What profit can a rational person with $\$M$ make conducting the following trade scheme once? He starts by using his $\$M$ money to buy[10] any bundle he wants from the consumer at her perceived prices p^c, then trades that bundle in any way he wants on the market given the true prices p, and finally resells whatever goods he acquired this way back to the consumer at her perceived prices p^c. To be clear, this scheme is conducted before the consumer decides on her consumption plan; she is not yet maximizing her preference, only accepting trades that give her a higher perceived budget for doing so, which is always desirable. For simplicity, our tie-breaking rule is that the consumer accepts trades which leave her budget unchanged.[11] To summarize, the third party will maximize $p^c \cdot y$ over all bundles y such that $p \cdot y \leq p \cdot x$, for some bundle x such that $p^c \cdot x \leq M$.

To solve this optimization problem, notice that the solution will have the bundle x maximize $p \cdot x$ over the set of bundles x such that $p^c \cdot x \leq M$. Indeed, making $p \cdot x$ larger increases the set of bundles y such that $p \cdot y \leq p \cdot x$. Because the objective function $p \cdot x$ is linear in x, an optimal solution to this problem is to spend the $\$M$ on a good ℓ with the highest price ratio when comparing true prices to perceived prices: $p_\ell/p^c_\ell \geq p_k/p^c_k$ for each good k (the computation is analogous to that when maximizing perfect-substitutes preferences). It remains to find a y that maximizes $p^c \cdot y$ under the constraint $p \cdot y \leq (p_\ell M)/p^c_\ell$. Similar reasoning reveals that an optimal solution is to spend the $\$(p_\ell M)/p^c_\ell$ on a good ℓ' with the highest price ratio when comparing perceived prices to true prices: $p^c_{\ell'}/p_{\ell'} \geq p^c_k/p_k$ for each good k. The profit in that case is $\$(p^c_{\ell'}/p_{\ell'})(p_\ell M)/p^c_\ell$. To summarize, the third party's maximal profit is:

$$\max_{\ell,\ell'} \frac{p_\ell/p_{\ell'}}{p^c_\ell/p^c_{\ell'}} M = \delta(p, p^c) M.$$

Thus, using δ to measure how far apart the true price vector is from the perceived one also determines the money pump multiplier for a rational party conducting the

[10]The consumer's endowment in goods is assumed to be large enough, or M is assumed to be small enough that the consumer can provides the good that the third party wants to buy.

[11]If she accepts only trades that strictly increase her budget, then the rational schemer can get as close as desired from the optimal profit calculated in the next paragraph, by leaving a little bit of surplus to the consumer in both trades involved in the money pump scheme.

simple trading scheme described above. In this view, (3b) amounts to placing an upper bound ($\frac{1}{1-\varepsilon}$) on this money pump multiplier.

2.3 Misperceived Tastes

Suppose instead that the consumer is subject to errors in assessing her utility tradeoffs, or that the modeler's data is missing contextual information affecting these tradeoffs. As will become clear, contrary to Section 2.2, restricting attention to convex preferences is important for this interpretation of ε-rationalizability.

How do consumers explore their budget sets? A rational consumer might contemplate all bundles at once, to find the best one every time she has to make a choice. Alternatively, to save on contemplation and thinking costs, she may test any given bundle by checking that there is no preferable alternative in its vicinity, thereby reaching her choice – an affordable bundle that has no better alternative in its vicinity – by tatonnement. Reaching choices through a series of small adjustments might be a better description of a consumer's thought process in some circumstances. It is unlikely that every time some prices change at the grocery store, a consumer reassesses the entire set of affordable bundles, carefully introspects about her preference relation over those bundles, and directly selects the bundle maximizing it globally. With convex preferences, rational choices would in fact arise from thinking locally and in steps, as above. If the consumer is boundedly rational, on the other hand, she may make mistakes when assessing local utility tradeoffs by introspection, using utility tradeoffs that are not quite compatible with her true ones. Alternatively, she may assess her tradeoffs accurately, but these tradeoffs vary to some extent with factors unobserved by the modeler (e.g., ripeness of the fruit, expiration dates). Using the notion uncovered in Section 2.1 to bound the distance between perceived and true utility gradients, or between the actual and modeled utility gradients, again leads to ε-rationalizability.

Looking at how our notion of misperceived tastes operates in specific classes of utility functions provides another motivation for ε-rationalizability. Consider, for instance, Cobb-Douglas preferences. If the consumer's true preference is captured by the utility function $\prod x_\ell^{\alpha_\ell}$, then a natural way to capture misperceived tastes or parameter misspecification on the part of the modeler is that the consumer maximizes $\prod x_\ell^{\beta_\ell}$ where the vector of exponents β may vary, but cannot be too far apart from α. This operationalizes, in a different setting, Rubinstein and Salant (2012)'s notion of a decision maker who uses only preferences that are 'close' to her true one. As the next proposition shows, ε-rationalizability is equivalent to this approach, not only for the Cobb-Douglas model,[12] but also for *any* additively separable preference. Formally, say u is *additively separable* if $u(x) = \sum_\ell u_\ell(x_\ell)$ for some strictly concave,[13] continuous and strictly monotone utility functions $u_\ell : \mathbb{R}_+ \to \mathbb{R}$.

Proposition 3 *For each $\ell = 1, \ldots, L$, let u_ℓ be a utility function over good ℓ. Then $\mathcal{D}$ is ε-rationalizable with respect to the additively separable utility $u(\cdot) = \sum_{\ell=1}^L u_\ell(\cdot)$ if and only if there is $\beta : \mathcal{D} \to \mathbb{R}_{++}$ such that for all $(p, x) \in \mathcal{D}$,*

$$x = \underset{\{p \cdot y \leq p \cdot x\}}{\operatorname{argmax}} \sum_{\ell=1}^{L} \beta_\ell(p, x) u_\ell(y_\ell), \textit{ and} \tag{4a}$$

$$1 - \varepsilon \leq \frac{\beta_\ell(p, x)}{\beta_{\ell'}(p, x)} \leq \frac{1}{1 - \varepsilon}, \textit{ for all } \ell, \ell' \in \{1, \ldots, L\}. \tag{4b}$$

The inequalities in (4b) simply state that the vector β of modified coefficients is not too far from the original unit vector of coefficients associated to u, using the notion uncovered in Section 2.1. In the Cobb-Douglas example, $u_\ell(\cdot) = \alpha_\ell \log(\cdot)$ for each ℓ. For intertemporal choices, with x_ℓ representing the amount of good ℓ in time-period ℓ, exponential discounting corresponds to the case $u_\ell(\cdot) = \delta^\ell \tilde{u}(\cdot)$ for some time-independent utility function $\tilde{u}$ (i.e., independent of ℓ). In that context, Proposition 3 implies all the consumer's errors can be attributed to misperception of the discounting function. Similarly, in a setting with risk, where each ℓ is a state of the world, all errors could be attributed to misperception of probabilities.

3. Testable Implications

The seminal work of Afriat (1967) shows how the 'Generalized Axiom of Revealed Preference' (GARP) captures the empirical content of rational choice. Formally, demand data is consistent with the maximization of some strongly monotone, continuous utility function if, and only if, it satisfies GARP. Given the widespread use of regularity (and differentiable, single-valued demand functions), subsequent work extended

[12] Technically, the proposition applies to Cobb-Douglas only when restricting attention to strictly positive bundles, as $log(x_\ell) \in \mathbb{R}$ only if $x_\ell > 0$.

[13] Strict concavity may seem much stronger than our usual requirement of strict quasi-concavity. However, they are almost the same in this additive setting: in a classic result which builds on Arrow's earlier observation, Debreu and Koopmans (1982) show that quasi-concavity of a continuous, additively separable utility function implies that all but one u_ℓ's must be concave, and the last must have features of concavity too.

Afriat's approach in that direction. Under regularity, consistency is equivalent to the classic 'Strong Axiom of Revealed Preference' (SARP).[14]

We now show how to build on these approaches to capture the empirical content of ε-rationalizability. To do this, we apply the methodology of de Clippel and Rozen (2018). The first step is to assume the consumer follows the theory, and identify necessary restrictions that the demand data reveals about her preference. While the consumer's preference is defined over the entire space of goods $\mathbb{R}^L_+$, it turns out that the only relevant restrictions apply to her preferences over the subset of bundles $X = \{x \in \mathbb{R}^L_+ \mid (p, x) \in \mathcal{D} \text{ for some } p \in \mathbb{R}^L_{++}\}$ that have been observed chosen.

Suppose the DM's preference over bundles is represented by the regular utility function u. For any observation $(p, x) \in \mathcal{D}$, ε-Rationalizability requires that there is a vector $g \in \partial u(x)$ such that:

$$\frac{p_\ell}{p_{\ell'}}(1-\epsilon) \leq \frac{g_\ell}{g_{\ell'}} \leq \frac{p_\ell}{p_{\ell'}}\left(\frac{1}{1-\epsilon}\right), \quad \forall\, \ell \neq \ell', \tag{5}$$

The linear inequalities in (5) thus define a convex cone $C_\varepsilon(p, x) \subseteq \mathbb{R}^L_{++}$ which must contain a quasi-gradient of u at the bundle x. By strict quasi-concavity of u, we know that $u(x) > u(x')$ for any bundle x' such that $g \cdot x' \leq g \cdot x$. Graphically, x is better than the alternatives below the separating hyperplane through x which is determined by the vector g. This preference comparison remains true when restricted to those bundles which have been observed chosen. That is, x is strictly preferred to all bundles in the set:

$$\Gamma(g, x) = \{x' \in X \setminus \{x\} \mid g \cdot x' \leq g \cdot x\}\}. \tag{6}$$

The difficulty, of course, is that the modeler does not know the consumer's utility function u, let alone which quasi-gradient applies from each $C_\varepsilon(p, x)$. Rather, the modeler wishes to test whether the data is ε-rationalizable in the first place. The observations above shed light on this. Let us presume that the data is ε-rationalizable and consider a datapoint $(p, x) \in \mathcal{D}$. There must exist $g \in C_\varepsilon(p, x)$ such that x is strictly preferred to all the elements in $\Gamma(g, x)$. In the language of de Clippel and Rozen (2018), this is a *lower-contour set restriction* on the consumer's preference $\succ$ over the set X of demanded bundles, as some set from the family $\{\Gamma(g, x) \mid g \in C_\varepsilon(p, x)\}$ must

[14]See Chiappori and Rochet (1987), Matzkin and Richter (1991), and Lee and Wong (2005).

be contained in the $\succ$-lower contour set of x. There is one such lower-contour set restriction for *each* $(p, x) \in \mathcal{D}$, generating a collection of lower-contour set restrictions $\mathcal{R}_\varepsilon(\mathcal{D})$ over the consumer's possible preference over demanded bundles.[15] Clearly, a necessary condition for ε-Rationalizability of the demand data $\mathcal{D}$ requires the collection of restrictions $\mathcal{R}_\varepsilon(\mathcal{D})$ to be *acyclically satisfiable*: there must exist a (strict) acyclic relation on X that simultaneously satisfies them. For instance, the consumer's true utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$, restricted to X, defines such an acyclic relation.

We have not yet shown, however, that acyclic satisfiability of $\mathcal{R}_\varepsilon(\mathcal{D})$ implies the existence of a regular utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$ that ε-rationalizes the data. The following result confirms this.

Proposition 4 *$\mathcal{D}$ is ε-rationalizable if and only if $\mathcal{R}_\varepsilon(\mathcal{D})$ is acyclically satisfiable.*

Proposition 4 shows that testing whether demand data $\mathcal{D}$ is ε-rationalizable amounts to checking whether the collection of restrictions $\mathcal{R}_\varepsilon(\mathcal{D})$ is acyclically satisfiable. How does one check this? de Clippel and Rozen (2018) shows that in the presence of lower-contour set restrictions, one checks acyclic satisfiability in the same way one checks SARP: by iteratively looking for candidate minimal elements of the unknown preference ordering over X. The key step is this: having determined thus far that the elements in $S \subset X$ can be ranked at the bottom of the ordering, is there any element in the set $X \setminus S$ of remaining, unranked alternatives that qualifies as a candidate-minimal element among those? An option x qualifies if and only if no other remaining elements fall below the separating hyperplanes going through x, since none of these elements need to be ranked below it. Formally, x qualifies if for every price vector p such that $(p, x) \in \mathcal{D}$, there exists $g \in C_\varepsilon(p, x)$ such that $\Gamma(g, x) \cap (X \setminus S) = \emptyset$. This step is easy to determine by linear programming.[16] Acyclic satisfiability holds if and only if we can identify a candidate minimal element in each iteration, thus constructing a complete ordering over the entire set of chosen bundles X. Importantly, the success or failure of this procedure is path independent: it does not depend on which candidate minimal element one picks in a step when multiple alternatives qualify.

[15] Notice that the arguments here accommodate the possibility that the same bundle x is chosen at two or more different price vectors: in such a case, there are simply multiple lower-contour set restrictions associated with the same bundle x.

[16] Developing this, x qualifies if and only if for each $(p, x) \in \mathcal{D}$ there is a vector g that satisfies the weak inequalities in (5) and the strict inequalities $g \cdot (x' - x) > 0$ for all $x' \in X \setminus S$ other than x. It is not difficult to transform this problem into one with only weak inequalities, for which linear programming techniques then apply.

While the formalization above may be notationally heavy, the concept is actually very simple, and a more colloquial description would be worthwhile. Indeed, $g \in C_\varepsilon(p, x)$ just means that g perturbs the price vector p, making it at most ε away according to our measure δ from Section 2. Graphically, one can think of g as tilting the true budget line while still going through the demanded bundle. We begin the test by seeking a candidate for worst overall element. An option x qualifies if, for every price vector p for which it is demanded, it is possible to tilt the p-budget line going through x (with the tilted price at most ε away from the true one) in a way that makes all other demanded bundles unaffordable. Once we have found a candidate x for worst-overall element, we can effectively erase it and move on, seeking the next-worst element x' using the same budget-tilting procedure to ensure that the remaining demanded bundles are not affordable. This process continues until there are no demanded bundles left, in which case the data is ε-rationalizable; or until we find ourselves stuck, in which case it is not. The idea of the proof is to use the tilted prices g as utility quasi-gradients of the ε-rationalizing utility function.

Our testing methodology for ε-rationalizability rests on the fact that the possible quasi-gradients are restricted to a set. The approach is portable to other situations where this feature arises. It is easily tweaked, for instance, to test ε-rationalizability with an added restriction of differentiability on the utility function,[17] but works beyond the context of ε-rationalizability as well. Consider a modeler who theorizes that a consumer's misperceived price for each good ℓ is some mixture of a default price p^d_ℓ and the true price: $p^c_\ell = p^d_\ell + m_\ell(p_\ell - p^d_\ell)$, where $m_\ell \in [0, 1]$ is the extent to which the consumer shifts the perceived price of the good from the default price to the actual one. This is a weaker restriction on perception than Gabaix (2014)'s sparse-max theory, which determines the vector $m = (m_1, \ldots, m_L)$ endogenously through a sparsity-based optimization using a concave utility function.[18] To test the generalization where any $m \in [0, 1]^L$ is conceivable, which yields predictions less precise than Gabaix's, but robust to a variety of theories of subjective price formation, one replaces $C_\varepsilon(p, x)$ with the convex set of possible gradients $C(p, p^d) = \{(p^d_\ell + m_\ell[p_\ell - p^d_\ell])^L_{\ell=1} \mid m \in [0, 1]^L\}$,

[17] With differentiability, $\partial u(x)$ is a singleton (up to positive rescaling). Hence for each $x \in X$, there must be $g \in C_\varepsilon(x) := \bigcap_{(p,x)\in\mathcal{D}} C_\varepsilon(p, x)$ such that x is strictly preferred to $\Gamma(g, x)$.

[18] His theory captures a decision maker who, for instance, realizes that spending time understanding the state of the Amazonian forest or interest rates in some distant country will be costly and yet have very little effect on her decision. The decision maker knows default values of such variables (e.g., long-run averages) and optimally allocates effort in determining what price estimate, somewhere between the default value and the true value, to use.

and proceeds with the analogous lower-contour set restrictions.

4. FOC-Departure Index

Demand data is either consistent with rationality, or not. In the imperfect world of actual data, it is more useful to have ways to quantify the degree to which data complies with a theory. ε-Rationalizability naturally lends itself to such measurements.

Definition 2 (FOC-Departure Index) *The* FOC-Departure Index (FDI) *is the infimum over all ε such that the data is ε-rationalizable.*[19]

For instance, the FDI of demand data arising from the maximization of a regular preference is zero. The FDI provides an alternative measure of goodness-of-fit to Afriat (1973)'s Critical Cost Efficiency Index (CCEI), which computes the largest percentage of the consumer's budgets that can be retained while eliminating all revealed-preference cycles (see also Varian (1990)). To formalize Afriat's index, define for $\sigma \in [0,1]$ a strict revealed preference $x \succ_{A,\sigma} y$ if $(p,x) \in \mathcal{D}$ and $\sigma p \cdot x > p \cdot y$; and a weak revealed preference $x \succeq_{A,\sigma} y$ if $\mathcal{D}$ contains a sequence of observations $(p^1,x^1),\ldots,(p^n,x^n)$ where $x^1 = x$, $x^n = y$, and for each $i \in \{1,\ldots,n-1\}$, either $x^i = x^{i+1}$ or $\sigma p^i \cdot x^i \geq p^i \cdot x^{i+1}$. Then the CCEI is the supremum of $\sigma \in [0,1]$ such that $x \succeq_{A,\sigma} y$ implies not $y \succ_{A,\sigma} x$ for all observed choices x, y.[20] While both the CCEI and FDI yield values between zero and one, the directions are reversed: the CCEI measures rationality, and is 1 when choices are rational, while the FDI measures departure from rationality, and is 0 when choices are rational. However, one can directly compare the FDI to Afriat's *in*efficiency measure $1-$CCEI, which is the minimal factor by which the consumer's budgets must be *shrunk* to eliminate revealed-preference cycles.

We start in Section 4.1 by illustrating the FDI in the simpler case of two commodities (as in most experiments on the subject). In addition to providing further

[19] Definition 1 remains meaningful when relaxing the notion of regularity in ε-rationalizability to allow *weak* quasi-concavity (it suffices for local optima to be global optima). We focused on *strict* quasi-concavity, as SARP is a bit easier to describe and work with than GARP. This has no impact on the value of the FDI: if $\mathcal{D}$ is ε-rationalizable when allowing quasi-concavity to be weak, then $\mathcal{D}$ is ε'-rationalizable when requiring quasi-concavity to be strict, for *all* $\varepsilon' > \varepsilon$.

[20] Varian (1990) generalizes Afriat's index to allow the (proportional) budget adjustment to vary per observation. Houtman and Maks (1985) propose another classic index, the smallest subset of the data that needs to be dropped to make it rationalizable. Halevy et al. (2018) points out that these indices can be viewed as being in the same umbrella class, with different restrictions on the possible budget adjustments, and different ways of aggregating across observations. We refer the reader to Halevy et al. (2018, Appendix B) for a thorough discussion of some other metrics that don't fall under this umbrella, such as Echenique, Lee and Shum (2011), Apesteguia and Ballester (2015), and Dean and Martin (2016).

intuition, these results will prove useful throughout this section and the next, providing a simple, closed-form formula for the FDI of demand data with two commodities.[21] Section 4.2 shows that Afriat's inefficiency measure is always smaller or equal to the FDI (independently of the number of commodities). We show by example that no systematic relation exists when taking power into account: demand data may appear closer to rationality under the CCEI than under the FDI, and vice versa. In Section 4.3, we revisit data from Choi et al. (2014) to illustrate how the FDI can easily be used in practice.

4.1 The Case of Two Commodities

We assume $L = 2$ throughout this subsection. As a start, consider demand data $\mathcal{D} = \{(p, x), (p', x')\}$ comprising only two observations. The next result provides a simple formula for computing the FDI of any such $\mathcal{D}$. Notice that with only two observations, a violation of SARP is a situation in which x, x' differ and each of these bundles is affordable when the other is chosen (i.e. $p \cdot x' \leq p \cdot x$ and $p' \cdot x \leq p' \cdot x'$).

Proposition 5 *Take $\mathcal{D} = \{(p, x), (p', x')\}$, and assume $p'_1/p'_2 \geq p_1/p_2$ without loss of generality. Then,*

$$\mathrm{FDI}(\mathcal{D}) = \begin{cases} \min\{\frac{p\cdot(x-x')}{p_2(x_2-x'_2)}, \frac{p'\cdot(x'-x)}{p'_1(x'_1-x_1)}\} & \textit{if } \mathcal{D} \textit{ violates SARP;} \\ 0 & \textit{otherwise.} \end{cases}$$

The proof of Proposition 5 reveals what underlies this formula. In particular, in the event of a SARP violation, the FDI directly relates to how far apart each of the two price vectors is from the vector $o(x, x')$ that is orthogonal to the line passing through x and x' (see Figure 1). Indeed, the FDI is the minimum between $f^{-1}(\delta(p, o(x, x')))$ and $f^{-1}(\delta(p', o(x, x')))$, where f is the function that associates to each $\varepsilon \in (0, 1)$ the fraction $1/(1-\varepsilon)$ appearing in the definition of ε-rationalizability.

The next result shows that, quite remarkably, the FDI of demand data with two commodities is simply the maximum of the FDI over all *pairs* of observations. Combined with Proposition 5, we have thus found a formula to compute the FDI for any number of observations.

[21] While we haven't found an analogous formula with three commodities or more, it takes only n applications of the enumeration procedure to identify the index with $\pm\frac{1}{2^n}$ precision. First, test for 1/2-rationalizability; then test 1/4-rationalizability if the previous test succeeds, and 3/4-rationalizability otherwise; and continue this recursively $n-2$ more times.

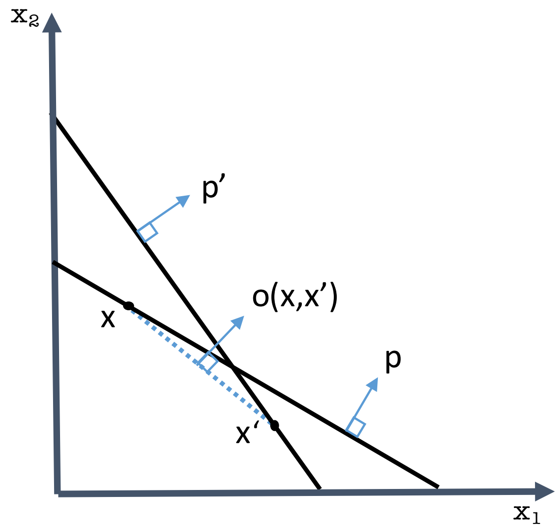

Figure 1: The construction of $o(x, x')$ for proving Proposition 5.

Proposition 6 *For $L = 2$, the FDI of any demand data $\mathcal{D}$ is the maximum of the FDI's (computed in Proposition 5) associated to each pair of observations in $\mathcal{D}$.*

Demand data is ε-rationalizable if and only if for every $x \in X$, there is $g(p, x) \in C_\varepsilon(p, x)$ such that the auxiliary demand data $\mathcal{D}' = (g(p, x), x)_{(p,x)\in\mathcal{D}}$ is rationalizable (in the classic sense) by a regular utility function u. That, in turn, is equivalent to $\mathcal{D}'$ satisfying SARP. One may conjecture then that Proposition 6 follows at once from Rose (1958).[22] That would be the case if we knew that for every $(p, x) \in \mathcal{D}$, there is $g(p, x) \in C_\varepsilon(p, x)$ such that for all $(p', x'), (p'', x'') \in \mathcal{D}$ the auxiliary demand data $\{(g(p', x'), x'), (g(p'', x''), x'')\}$ satisfies SARP. However, multiple SARP violations may occur when pairing a given observation (p, x) with different other observations, and one may have to consider different vectors $g(p, x)$ to achieve ε-rationalizability on those pairs.[23] This makes the proof significantly harder.

[22]Rose observes that, with only two commodities, satisfying SARP over pairs of observations guarantees satisfying SARP over sequences of observations of any length.

[23]For instance, suppose $p = (1, 1)$, $p' = (1, 2)$, $p'' = (2, 1)$, $x = (10, 10)$, $x' = (1, 18)$, and $x'' = (18, 1)$. Both $\{(p, x), (p', x')\}$ and $\{(p, x), (p'', x'')\}$ are ε-rationalizable with $\varepsilon = 1/7$. However, the p' budget line cannot be tilted to eliminate the revealed preference of x' over x, while keeping $\delta(p', g(p', x'))$ below $1/(1-\varepsilon)$. Instead, one must make the p-budget line flatter (e.g. taking $g(p, x) = (5, 6)$ guarantees SARP while keeping $\delta(p, g(p, x)) < 1/(1-\varepsilon)$). Thus, for $\{(g(p, x), x), (g(p', x'), x')\}$ to satisfy SARP, $g(p, x)$ must put significantly more weight on good 2 than good 1. A similar argument for $\{(p, x), (p'', x'')\}$ implies that a different $g(p, x)$, one that puts significantly more weight on good 1 than good 2, must be considered to get $\{(g(p, x), x), (g(p'', x''), x'')\}$ to satisfy SARP.

4.2 Relation to the CCEI

From now on, we return to the general case with no restriction on the number of commodities. As the CCEI and FDI pertain to very different variables – adjusting incomes for the CCEI, versus adjusting tradeoffs for the FDI – one might think these measures are unrelated. However, we show that there is a clear relationship between the two. In particular, a small departure from first-order conditions (in the sense of our measure) means only small budgetary adjustments are needed to eliminate revealed-preference cycles.

Proposition 7 *For any demand data $\mathcal{D}$,* $1 - \mathrm{CCEI}(\mathcal{D}) \leq \mathrm{FDI}(\mathcal{D})$.

As will be illustrated in our application to portfolio-choice data in Section 4.3, the inequality in Proposition 7 can hold strictly for some datasets, and with equality for others. Moreover, the measures can offer opposing assessments of different datasets: in some cases, the CCEI might deem one dataset closer to rationality than another, while the FDI determines the opposite.

Of course, whether consistency with rationality is remarkable depends, at least to some extent, on the combination of budget sets being tested. For instance, rationality is impossible to refute when all budget sets are related by inclusion. By contrast, a SARP violation becomes possible in intersecting budget sets that are not related by inclusion. In that sense, the specific value of the CCEI or the FDI derived from a consumer's actual choice is not that informative without being contrasted against the distribution of those indices arising under some alternative behavioral hypothesis.

While different criteria have been proposed over the years to capture power, we focus here on the approach that is most often applied in experimental papers. This method, suggested by Bronars (1987) and inspired by Becker (1962), proposes to use as a reference point the distribution of CCEI's arising from a random collection of choices, under the assumption that each bundle on the frontier of a budget set is equally likely. Most experimental papers argue that the rational choice model captures observed choices rather well, because the distribution of CCEIs arising from the data is a significant FOSD shift towards lower values of Afriat's index.

Such an approach can be replicated using the FDI instead. Interestingly, while we know from Proposition 7 that $1 - \mathrm{CCEI} \leq \mathrm{FDI}$, this does not mean that subjects will necessarily appear less rational when applying Bronars' methodology to the FDI. Indeed, the distribution of FDI's for the randomly generated demand data used as the reference point will itself shift towards higher values. Here is a simple theoretical example to illustrate this point.

Example 1 *Consider demand data $\mathcal{D} = \{(p, x), (p', x')\}$ with two observations and a SARP violation, where the price vectors are $p = (1, 2)$ and $p' = (2, 1)$ and the demanded bundles are $x = (3, 6)$ and $x' = (6, 3)$. We first observe that the largest possible FDI, for any pairs of choices on the boundary of these two budget sets, is $1/2$. To see this, let y, y' be two distinct bundles such that y is on the p-budget line, y' is on the p'-budget line, and there is a SARP violation: $p \cdot y' \leq p \cdot y = p \cdot x$ and $p' \cdot y \leq p' \cdot y' = p' \cdot x'$. By Proposition 5, the FDI of such choices is*

$$1 - \frac{1}{2} \max\{\gamma, \frac{1}{\gamma}\}, \text{ with } \gamma = \frac{y'_1 - y_1}{y_2 - y'_2}.$$

Hence the largest FDI is reached with $\gamma = 1$, which occurs with probability zero when drawing uniformly from bundles on the boundary of the two budget sets. Hence, the probability that the FDI of a randomly drawn demand data is larger or equal to that of the demand data above, which achieves that maximum FDI, is zero. By contrast, there is a strictly positive probability that the Afriat inefficiency of a randomly drawn demand data will be bigger than that of $\mathcal{D}$.

For an example where the comparison is opposite, consider the demand data $\mathcal{D}' = \{(p, x), (p', z)\}$ where p, p', x are as above, but z amounts to spending the budget under p' entirely on good 1, i.e., $z = (7.5, 0)$. Notice that the Afriat inefficiency of demand data picked on the budget lines associated to p and p' is smaller or equal to that of $\mathcal{D}'$ if and only if the bundle picked on the p-budget line is strictly to the left of x and the bundle picked on the p'-budget line is strictly to the right of x' (defined above). It is easy to check that any such demand data also has a FDI larger than that of $\mathcal{D}'$, but also that there is a positive mass of other bundle combinations leading to a larger FDI than that of $\mathcal{D}'$.[24] Thus, this time, the probability that the FDI of a randomly drawn demand data is larger than that of $\mathcal{D}'$, is larger than the probability that the Afriat inefficiency of a randomly drawn demand data is larger than that of $\mathcal{D}'$.

[24] For instance, for all $\hat{z}$ halfway between x' and z on the p'-budget line, there is a positive mass of bundles $\hat{x}$ to the right of x on the p-budget line for which the FDI of $\{(p, \hat{x}), (p', \hat{z})\}$ is strictly larger than that of $\mathcal{D}'$.

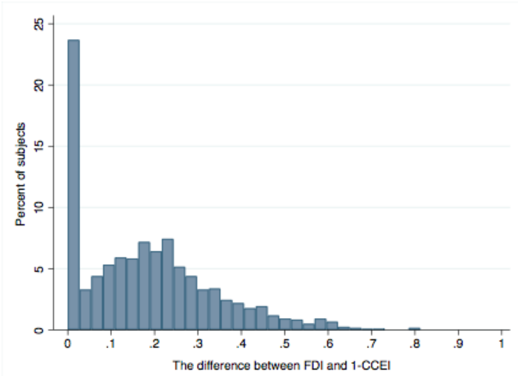

(a) Illustration of Proposition 7.

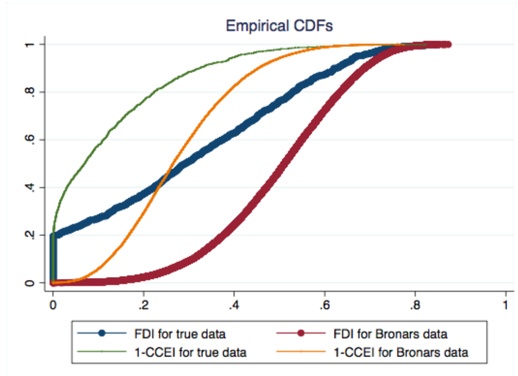

(b) Taking power into account.

Figure 2: Illustration using data from Choi et al. (2014).

4.3 Illustration Using Data

To illustrate, we present our methodology in the simpler setting of recent lab experiments on risky portfolio choices. Given two equally-likely states, subjects face multiple linear budget sets and decide how to allocate money across states given these constraints. There are thus two commodities: money in state 1, and money in state 2. This is the setting of Choi et al. (2014), whose demand data we revisit through the lens of our theoretical results.

We compute the FDI using the simple formula provided by Proposition 6, as $L = 2$. Proposition 7 tells us that $1 - \text{CCEI} \leq \text{FDI}$, and the histogram of $\text{FDI} - (1 - \text{CCEI})$ in Figure 2(a) reveals the extent by which these two measures actually differ in the data. For just under one-quarter of subjects, the $\text{FDI} = 1 - \text{CCEI}$; among those for whom these differ, the modal difference is around 0.2. As they both measure departure from rationality, the FDI and $1 - \text{CCEI}$ are correlated: the Spearman correlation is 0.8481, and the null hypothesis that the two are independent is strongly rejected (p-value of 0.0000). One index is not simply a monotone transformation of the other. In fact, for around 15% of all pairs of subjects in the experiment, the CCEI and the FDI offer opposite rankings of departure from rationality: that is, Anne is considered more rational than Bob under one index, but Bob is considered more rational than Anne under the other.

To take power into account, we perform a Bronars exercise by repeatedly drawing random choices for each sequence of 25 budget sets tested in Choi et al. (2014), for a total of 23,640 'random consumers' making 25 choices each. We compute the associated measures FDI and $1 - \text{CCEI}$ for the Bronars data. The empirical CDFs for both the true and Bronars data are plotted in Figure 2(b). Consistent with Proposition 7, for both the true and Bronars data we see the empirical CDFs first-order stochastically improve when moving from $1 - \text{CCEI}$ to the FDI. Thus, following our discussion of power in Section 4.3, data need not appear less rational under the FDI than the CCEI. As it turns out, both the FDI and $1 - \text{CCEI}$ exhibit a large, first-order stochastic shift down when moving from the true data to the Bronars data. This confirms the robustness of Choi et al. (2014)'s preliminary finding that there is a significant amount of rationality in their data.

5. Further Restrictions on Preferences

In many contexts, one is willing to impose further properties on the utility function beyond regularity. Classic examples in consumer theory include quasi-linearity, homotheticity and additive separability. Expected utility is a typical assumption for choice under risk, as are exponential or hyperbolic discounting for time preferences.

Let $\mathcal{U}$ be the subclass of regular utility functions under consideration. Demand data is *ε-rationalizable with respect to $\mathcal{U}$* if it satisfies Definition 1 with the added requirement that u belongs to the class $\mathcal{U}$. Given a dataset, one can then define its FDI *with respect to $\mathcal{U}$*, denoted $\text{FDI}_{\mathcal{U}}$, as the infimum over ε such that the data is ε-rationalizable with respect to $\mathcal{U}$. Of course, the FDI with respect to $\mathcal{U}$ weakly increases as $\mathcal{U}$ becomes smaller.

In Section 5.1, we develop two important connections with Afriat's CCEI and tests of classic rationalizability in the spirit of Afriat's inequalities, that hold for any subclass of regular utility functions $\mathcal{U}$. In Section 5.2, we apply these ideas to recent experimental data pertaining to choices from budget sets involving risk.

5.1 Connections with Afriat and Afriat's Inequalities

As in Halevy et al. (2014), one can define the CCEI with respect to $\mathcal{U}$, denoted $\text{CCEI}_{\mathcal{U}}$, as the largest share of income that can be retained such that there exists a utility function $u \in \mathcal{U}$ for which each demanded is strictly preferred to all bundles in the associated shrunken budget set. Perhaps surprisingly, both the proof and statement of Proposition 7 *extend verbatim* when restricting to classes of utility functions.

Proposition 8 *For any $\mathcal{U}$ and demand data $\mathcal{D}$, $1 - \text{CCEI}_{\mathcal{U}}(\mathcal{D}) \leq \text{FDI}_{\mathcal{U}}(\mathcal{D})$.*

Proposition 8 provides a lower bound for $\mathrm{CCEI}_{\mathcal{U}}$, which may be of independent interest. While easily defined, figuring out how to compute it may be challenging. By contrast, classic tests of rationalizability extend at once to compute the $\mathrm{FDI}_{\mathcal{U}}$. Indeed, as should be clear from its definition, ε-rationalizability with respect to $\mathcal{U}$ is equivalent to finding price vectors $(p^c(p,x))_{(p,x)\in\mathcal{D}}$ such that

$$(7) \qquad (1-\varepsilon)\frac{p_\ell}{p_{\ell'}} \leq \frac{p^c_\ell(p,x)}{p^c_{\ell'}(p,x)} \leq \frac{1}{1-\varepsilon}\frac{p_\ell}{p_{\ell'}}, \text{ for all } \ell,\ell' = 1,\ldots,L \text{ and } (p,x)\in\mathcal{D}$$

and the modified demand data $\mathcal{D}^c = \{(p^c(p,x),x)|(p,x)\in\mathcal{D}\}$ is rationalizable with respect to $\mathcal{U}$. At the same time, rationalizability with respect to standard classes of utility functions is often tested by checking whether certain sets of (oftentimes linear or bilinear) Afriat-like inequalities admit a solution.[25] Though they provide less insight into the theory than SARP-like characterizations (e.g., as developed in Section 3 for the general class of regular utility functions), such inequality-based tests are equally useful for testing purposes.

Combining these ideas, the next proposition explains how those classic tests can be repurposed for our intents. To state the result, note that any test of rationalizability with respect to $\mathcal{U}$ which is in the spirit of Afriat's inequalities can be described by a mapping $S_{\mathcal{U}}$, which associates to each demand data $\mathcal{D}$ a set of inequalities $S_{\mathcal{U}}(\mathcal{D})$ that has a solution if and only if $\mathcal{D}$ is rationalizable with respect to $\mathcal{U}$.

Proposition 9 *For any $\mathcal{U}$ and test of rationalizability $S_{\mathcal{U}}$, the demand data $\mathcal{D}$ is ε-rationalizable with respect to $\mathcal{U}$ if and only if there exists a solution to the system of inequalities listed in (7) and in $S_{\mathcal{U}}(\mathcal{D}^c)$, with an extra variable $p^c(p,x)\in\mathbb{R}^L_{++}$ for each $(p,x)\in\mathcal{D}$.*

This umbrella result is very convenient for computational purposes, as it applies to any subclass of regular preferences for which there is an existing result in the spirit of Afriat's inequalities. One simply checks a more permissive system of inequalities than required for the test of standard rationalizability with respect to $\mathcal{U}$: the true price vector in each observation (p,x) is replaced with a misperceived price vector $p^c(p,x)$, which must be close to p in the sense of our measure δ from Section 2 (itself a linear inequality). Of course, Proposition 9 also holds when $\mathcal{U}$ comprises the entire

[25]See Diewert (2012) among others, for an overview of classic tests of rationalizability with respect to homothetic or additively separability utility functions, and Brown and Calsamiglia (2007) for quasi-linear utility functions.

class of regular utility functions, in which case ε-rationalizability can be checked by the above tweak of the original Afriat (1967) inequalities. Conceptually, such a test is of a different nature than the one we develop in Proposition 4 of Section 3. Afriat (1967)'s characterization is multifaceted, offering several equivalent formulations. One formulation he offers is based on inequalities, which are somewhat difficult to interpret; another formulation he offers relates to revealed preference and is more insightful about the theory (Varian, 1982). Proposition 4, which is couched in terms of acyclic satisfiability of revealed preference restrictions, is in the latter vein.

5.2 Illustration with Portfolio Choice Data

In a contemporaneous and independent paper, Echenique et al. (2019) define and study what we would call ε-rationalizability for a risk-averse expected utility maximizer, except for the use of a different scale to measure departures from expected utility (the parameter e they use is equal to $\varepsilon/(1-\varepsilon)$).[26] In addition to their theoretical characterizations, they perform the important exercise of applying the new measure to a wide array of previously collected experimental data, which so far had been analyzed using the rational benchmark and/or a parametric approach. Our development of ε-rationalizability and the FDI for any class of regular utility functions allows us to add further layers of insight to some of the fundamental questions they study.

Echenique et al. empirically examine the relation between Afriat's CCEI and the new measure of goodness-of-fit for expected utility. Let $\mathcal{EU}^r$ denote the class of continuous, strictly monotone and strictly risk averse expected utility functions. Of course, as they point out, one expects subjects with a small $\text{FDI}_{\mathcal{EU}^r}$ to have a large CCEI, as being close to expected utility maximization implies a fortiori being almost rational. The relationship is more precise and general than this. As we showed, the FDI is larger or equal to $1 - \text{CCEI}$ (see Section 4), and $\text{FDI}_{\mathcal{EU}^r} \geq \text{FDI}$ since expected utility preferences are rational. Hence $\text{FDI}_{\mathcal{EU}^r} \geq 1 - \text{CCEI}$.[27]

By changing both the reference class of preferences and the way departures are measured, $\text{FDI}_{\mathcal{EU}^r}$ is conceptually two steps away from the CCEI. Using the FDI instead of the CCEI (which makes a difference, see Figure 2(a)) has the advantage of placing the spotlight on the dimension of interest: assessing how much more stringent expected utility is from rationality.

[26]They require only weak risk aversion; this has no effect on the FDI (see Footnote 19).
[27]Translating this inequality for their parametrization gives $\text{CCEI} \geq e/(1+e)$.

In fact, it is also important to assess how much more stringent expected utility is from the plain maximization of a utility function over *lotteries*. Notice that states have no meaning to subjects beyond the monetary amounts received. Though not required by rationality itself, subjects should view bundles simply as lotteries, which means their preferences should be *state independent*. Let $\mathcal{SI}$ denote the class of regular preferences that are state independent.[28]

To improve our understanding of $\text{FDI}_{\mathcal{EU}^r}$, we can decompose it as follows:[29]

$$\text{FDI}_{\mathcal{EU}^r} = \text{FDI} + [\text{FDI}_{\mathcal{SI}} - \text{FDI}] + [\text{FDI}_{\mathcal{EU}^r} - \text{FDI}_{\mathcal{SI}}].$$

Thus the degree of misspecification $\text{FDI}_{\mathcal{EU}^r} - \text{FDI}$ in using an expected utility preference is decomposed into the degree of misspecification from using a state independent preference, a very mild requirement that is common to all preferences recognizing that the bundles are lotteries, and the further misspecification from using the much more demanding expected utility form.

In the Appendix (see Proposition 11), we establish that for two equally likely states, demand data $\mathcal{D}$ is ε-rationalizable by a regular preference that is state independent if, and only if, the mirror-extended dataset

$$\overline{\mathcal{D}} = \mathcal{D} \cup \{((p_2, p_1), (x_2, x_1)) \mid (p, x) \in \mathcal{D}\} \tag{8}$$

is ε-rationalizable by a regular preference.[30] Thanks to this observation, we can easily compute $\text{FDI}_{\mathcal{SI}}$ by applying Proposition 6: it is the largest FDI (without requiring state independence) over all pairs of observations from the mirror-extended dataset.

---

[28]With equally likely states, state independence is equivalent to *symmetry*, or being indifferent between (x_1, x_2) and (x_2, x_1). With equally likely states and monotone preferences, state independence is also equivalent to *first-order stochastic monotonicity*, or preferring a lottery that is first-order stochastically superior to an alternative. If states were not equally likely, then symmetry would be unappealing, and first-order stochastic monotonicity would further restrict state independence by imposing an extra restriction on the preference over lotteries. The idea of approximately rationalizing the data with utility functions satisfying such properties was studied by Choi et al. (2014) using the CCEI. Echenique et al. (2019) show that there is a positive relationship between the frequency with which such properties are violated and their version of $\text{FDI}_{\mathcal{EU}^r}$.

[29]See Halevy et al. (2018) for similar decompositions with the CCEI for parametric classes.

[30]Necessity is obvious. Sufficiency is trickier. Suppose a regular u ε-rationalizes $\overline{\mathcal{D}}$. There is no reason to believe that u is state independent. One can easily modify u to symmetrize it (e.g., $\hat{u}(x) = u(x)$ if $x_2 \leq x_1$, and $= u(x_2, x_1)$ if $x_2 \geq x_1$), but the resulting preference is typically not convex anymore. Convexity would be preserved if any $g \in \partial u(x)$ with $x_2 < x_1$ is such that $g_1 < g_2$; that is, ensuring that below the diagonal, indifference curves are flatter than lines of slope -1. A

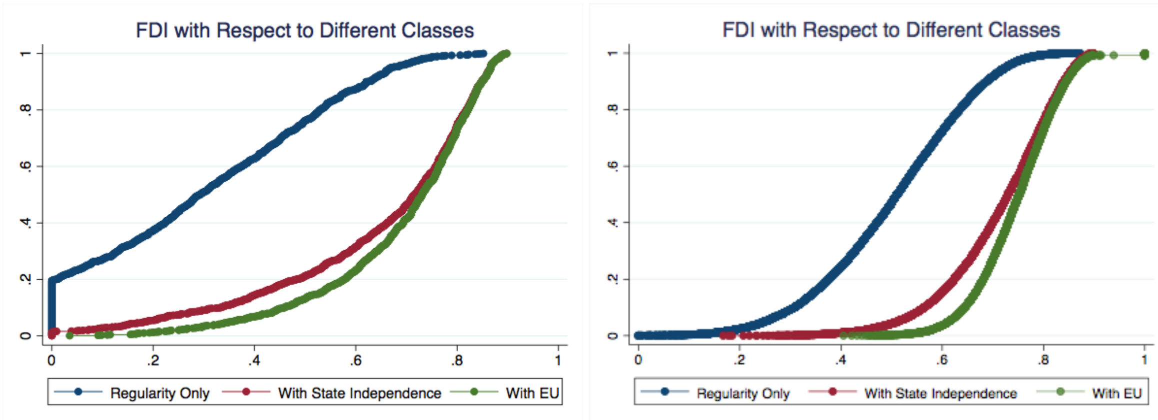

(a) For the true data.

(b) For the Bronars data.

Figure 3: Empirical CDF's of the FDI with respect to the class of regular utility functions and the subclasses of state-independent utility $\mathcal{SI}$ and risk-averse expected utility $\mathcal{EU}^r$.

Figure 3(a) depicts the empirical CDF's of the FDI, $\text{FDI}_{\mathcal{SI}}$, and $\text{FDI}_{\mathcal{EU}^r}$ for the data, where $\text{FDI}_{\mathcal{EU}^r}$ is computed using Proposition 9, by repurposing the classic test of Green and Srivavastra (1986) for risk-averse expected utility. It may come as a surprise that adding the mild requirement of state independence has a big impact, while the much more substantial restriction of adding expected utility on top of state independence has a much smaller impact. In fact, for roughly 65% of subjects in the actual experiment, we have that $\text{FDI}_{\mathcal{EU}^r} \approx \text{FDI}_{\mathcal{SI}} > \text{FDI}$. Consider the EU-misspecification ratio

$$\frac{\text{FDI}_{\mathcal{EU}^r} - \text{FDI}_{\mathcal{SI}}}{\text{FDI}_{\mathcal{EU}^r} - \text{FDI}}$$

across subjects, which is the proportion of misspecification of imposing expected utility instead of any regular preference over *bundles*, that is attributable to the misspecification for imposing expected utility instead of any regular preference over *lotteries*. Figure 4 suggests the empirical CDF of the misspecification ratio for the true data is not much better than that for the randomly-generated, Bronars data.[31]

This feature reflects the experiment's limited power to detect the relative validity of expected utility beyond plain preference maximization over lotteries. To gain some intuition, suppose a subject were to randomly choose a bundle on the frontier of a

key step of the proof is then to show the existence of such a utility function that ε-rationalizes $\overline{\mathcal{D}}$.

[31] For comparability, we use the same, realized Bronars dataset throughout.

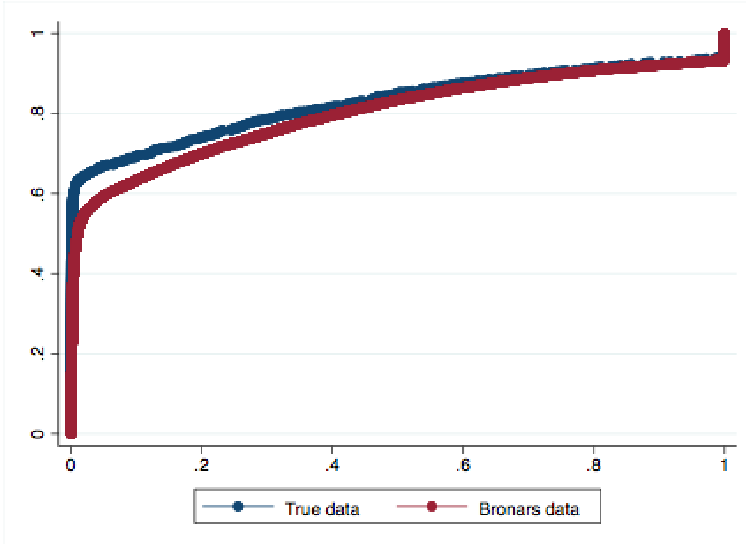

Figure 4: Empirical CDFs of the EU-misspecification ratio.

budget set $p_1 x_1 + p_2 x_2 = 1$, with $p_1 > p_2$. The probability that she appears rational ($\text{FDI} = 0$) is 1: there is too little information to refute rationality. However, $\text{FDI}_{\mathcal{SI}}$ is strictly positive with probability $p_2/(p_1 + p_2)$ (as any choice below the diagonal would violate it), and $\text{FDI}_{\mathcal{SI}} = \text{FDI}_{\mathcal{EU}^r}$ with probability 1. This is because rationalization by an expected utility preference comes for free once the price vector has been twisted enough to get rationalization by a state-independent preference.

Of course, the experiment tests subjects in many more budget sets, which makes the comparison more complex. Figure 3(b) depicts the same objects as in Figure 3(a), this time for the randomly-generated Bronars dataset. Though specific numbers change, of course, a similar general pattern arises: relatively speaking, imposing state independence is a larger leap from rationality, than expected utility is from state independence. This confirms our intuition that subjects would have to answer more complex questions (e.g., varying the probabilities of states) if we wish to gain a deeper understanding of the adequacy of expected utility beyond more general preferences over lotteries. This reinforces a related point made by Polisson, Quah and Renou (2019), who noticed (using the CCEI and a different dataset) that power to test expected utility conditional on rationalizability by a first-order stochastically monotone utility function can be very low.

6. Beyond Linear Budget Sets

In an empirical study of nonlinear pricing in electricity markets, Ito (2014) points out that while optimization requires understanding marginal prices, "nonlinear pricing and taxation complicate economic decisions by creating multiple marginal prices for the same good." Indeed, the empirical literature on non-linear pricing, with applications to labor and taxation (see the survey of Saez, Slemrod and Giertz (2012)) and utilities markets (see Ito (2014) and references therein) provides evidence that consumers respond to nonlinear pricing in a manner inconsistent with classic theories. These studies illustrate the importance of relaxing the assumption of perfect collinearity between marginal prices and the utility gradient.[32]

The analysis developed so far lends itself very well to such an endeavor. *Generalized demand data* $\mathcal{D}$ comprises a finite collection of pairs (P, x), where $P = (P_1, \ldots, P_L)$ is a collection of strictly increasing and differentiable price functions P_ℓ, and $x \in \mathbb{R}^L_{++}$ is the consumption bundle demanded for the budget set

$$B(P, x) = \{y \in \mathbb{R}^L_+ | \sum_\ell P_\ell(y_\ell) \leq \sum_\ell P_\ell(x_\ell)\}.$$

Indeed, $P_\ell(q)$ represents the total price to pay for buying a quantity $q \geq 0$ of good ℓ. Naturally, $P_\ell(0) = 0$ for each ℓ. We presume that $B(P, x)$ is convex (as would be true if P_ℓ is convex for all ℓ), in which case the first-order conditions are not only necessary but sufficient for an optimum.

Suppose for the moment that P is differentiable at the bundle x. Then, the *marginal price vector* at x is given by $(P_1'(x_1), P_2'(x_2), \ldots, P_L'(x_L))$, and the *marginal price ratio* associated to any two goods ℓ, ℓ' is the following strictly positive number:

$$MPR^P_{\ell,\ell'}(x) = \frac{P_\ell'(x_\ell)}{P_{\ell'}'(x_{\ell'})}.$$

More generally, the frontier of the budget set might be kinked. For instance, progressive taxation generates a convex budget set with a piecewise-linear frontier. Let $\partial P(x)$ be the set of strictly positive vectors defining the supporting hyperplanes of

[32]See Forges and Minelli (2009) for a study of rationalizability with generalized budget sets, and references therein.

$B(P,x)$ at x:

$$\partial P(x) = \{p \in \mathbb{R}^L_{++} \mid \forall y : \sum_{\ell} P_\ell(y_\ell) \geq P_\ell(x_\ell) \Rightarrow p \cdot y \geq p \cdot x\}.$$

Differentiability at x means $\partial P(x)$ is just the marginal price vector, up to rescaling.

Classic rationalizability (for a strictly convex preference) just means that the nonlinear budget set $B(P,x)$ and the upper-contour set of x intersect only at the bundle x. That is, there is some positive vector which defines a hyperplane through x that separates the two sets. Just as before, we can relax this condition using ε-rationalizability. In a fully differentiable world, such a relaxation means finding a regular utility function u such that $1-\varepsilon \leq MRS^u_{\ell\ell'}(x)/MPR^P_{\ell\ell'}(x) \leq 1/(1-\varepsilon)$ for each $\ell \neq \ell'$ and each $(P,x) \in \mathcal{D}$. More generally, the notion is formalized as follows.

Definition 3 (ε-Rationalizability) *For $\varepsilon \in (0,1)$, the generalized demand data $\mathcal{D}$ is* ε-rationalizable *if there exists a regular utility function u such that for all $(P,x) \in \mathcal{D}$, there are vectors $g \in \partial u(x)$ and $p \in \partial P(x)$ such that the same inequality* (2) *as before holds for each $\ell \neq \ell'$:*

$$1-\varepsilon < \frac{g_\ell/g_{\ell'}}{p_\ell/p_{\ell'}} < \frac{1}{1-\varepsilon}.$$

Then u is said to ε-rationalize the generalized demand data.

A couple of remarks are in order. First, we retain the name ε-rationalizability; this cannot create any confusion, since Definition 3 boils down to Definition 1 with linear prices. Second, recalling Section 2.1, the above inequalities simply ensure that we can always find a 'utility gradient' $g \in \partial u(x)$ and a 'marginal price vector' $p \in \partial P(x)$ that are not too far from each other: $\delta(g,p) \leq \frac{1}{1-\varepsilon}$, for all $(P,x) \in \mathcal{D}$. Lastly, the FDI readily extends to generalized demand data, as the infimum over all ε such that the generalized demand data is ε-rationalizable.

The next result follows at once from Definitions 1 and 2. It is powerful nevertheless, as it shows that all the results derived in Sections 3 and 4 can be leveraged for ε-rationalizability of generalized datasets.[33] We have thus found an easily computable

[33]Moreover, as we discussed at the end of Section 3, the testing approach is portable to some specific theories of misperception. For instance, in the context of nonlinear pricing, one might be interested in testing the theory of a consumer who maximizes a regular utility function u given a perceived price vector p^c that is a convex combination of average and marginal prices, and restricted to belong to $C(P,x) = \{[\alpha P'_\ell(x_\ell) + (1-\alpha)\frac{P_\ell(x_\ell)}{x_\ell}]_{\ell \in L} | \alpha \in [0,1]\}$.

notion of goodness-of-fit for consumer data involving nonlinear prices.

Proposition 10 *A generalized dataset $\mathcal{D}$ is ε-rationalizable if, and only if, there is a 'linearized' dataset $\mathcal{D}'$, given by a set of observations (p, x) with $p \in \partial P(x)$ for all $(P, x) \in \mathcal{D}$, which is ε-rationalizable.*

Whenever the demanded bundle x lies at a point of differentiability of the budget set $B(P, x)$, the linearized price p in Proposition 10 above is simply the marginal price vector $(P_1'(x_1), P_2'(x_2), \ldots, P_L'(x_L))$. Thus testing ε-rationalizability when P is differentiable can be done using our earlier tests. For more general pricing, testing may require considering more variables than before. So far we had to determine whether the demand data is consistent with a reference model, for some utility tradeoffs that are not too far from the known price ratios. Now, at kinks in the pricing schedule, we have to also guess what the marginal price ratios might be. Testing methods from Sections 3 and 4 extend though, as acceptable price vectors at x must belong to $\partial P(x)$, which is a convex cone. We skip the technical details for the general case, and instead illustrate the idea by means of an intuitive example.

Example 2 *Consider demand data over two goods: x_1 corresponds to time spent on leisure (as opposed to work), and x_2 corresponds to a composite consumption good. For simplicity, the price of good 2 is always normalized to 1, in which case the price of good 1 can be thought of as the real hourly wage. The consumer's endowment is the number H of waking hours (for instance, over the course of a year). She decides how many hours, given H, to work. There is a progressive tax system: $\mathbb{R}_+$ is partitioned in successive intervals I_1, I_2, $\ldots$, and the marginal tax rate in income bracket I_k is $t_k \in [0,1]$, with $t_1 \leq t_2 \leq t_3 \leq \ldots$. Thus for a total yearly wage $W \in I_k$, she will pay a total tax equal to $t_k(W - a_k) + \sum_{j=1}^{k-1} t_j(b_j - a_j)$, where $I_j = [a_j, b_j)$ for each j. Clearly, this leads to convex budget sets with piecewise-linear frontiers, which will change as the tax code and hourly wages vary.*

Suppose the consumer can allocate $4,000$ hours each year between work and leisure. She faces a different tax system on two different years, while her wage remains constant, at \$25 an hour. In year 1, income up to $\$25K$ is taxed at a 20% marginal rate, while the marginal tax rate for higher incomes is $33\frac{1}{3}\%$. In year 2, income up to $\$25K$ is not taxed, while higher incomes face a marginal tax rate of $66\frac{2}{3}\%$. Selected bundles are $x = (3K, 20K)$ for year 1, and $x' = (1.5K, 37.5K)$ for year 2. The situation is depicted in Figure 5.

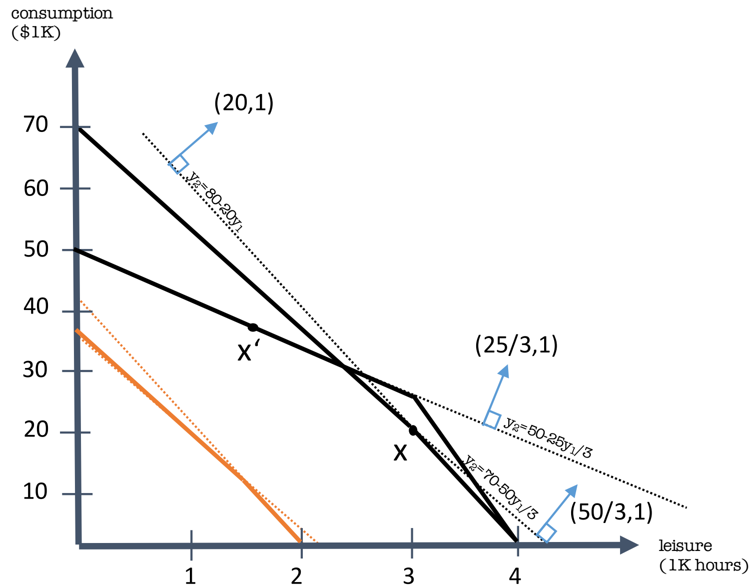

Figure 5: Illustration of the taxation example.

Clearly this demand data is not rationalizable, since x is revealed strictly preferred to x', and vice versa. What is the FDI? Since the budget frontier is differentiable at x', the demand data is ε-rationalizable if, and only if, the modified demand data, where the budget set on year 2 is expanded to $\{(y_1, y_2)|\frac{25}{3}y_1 + y_2 \leq 50,000\}$, is ε-rationalizable. We can also linearize the budget set on year 1, but here we are free to pick any price vector that is orthogonal to the budget set at x, that is, any convex combination of $(50/3, 1)$ and $(20, 1)$ (the normalized extreme vectors of that cone). Clearly, eliminating the revealed preference of x over x' by twisting the putative price as little as possible means that one should use $(50/3, 1)$, with the budget set on year 2 now expanded to $\{(y_1, y_2)|\frac{50}{3}y_1 + y_2 \leq 70,000\}$. Applying Proposition 5, the FDI of our original demand data is $5/17.5 \approx 28.5\%$.

As should be clear by now, the FDI provides a simple measure of rationality that is portable to all settings where the first-order approach applies. By contrast, the CCEI applies only if it is clear what it means to 'shrink' a budget set. That meaning becomes ambiguous beyond linear budget sets. One option is to pull budget frontiers proportionally towards zero. While seemingly close in spirit to Afriat's original idea, preserving the shape of budget sets may be at odds with the source of non-linearity. In the example above, shrinking the first budget set by one half (as depicted by the orange budget set) would implicitly mean that the consumer faces a marginal tax rate of 20% only on the first \$10K, which is not the case. Alternatively, one could shrink the available hours H which, in this case, amounts to shifting the budget frontier to the left instead of towards 0. While both methods are equivalent should prices be linear, which is the 'right' one to pick otherwise is unclear.

7. Beyond Convex Preferences

Quasi-concavity played a central role in our analysis. Definition 1 becomes trivially satisfied without it, and the associated FDI becomes uninformative (systematically equal to zero). Indeed, one can always draw indifference curves to match local restrictions at the finitely many bundles appearing in the demand data. As we all know, the first-order conditions provide little information unless accompanied by the second-order condition to guarantee a global optimum.

That being said, Proposition 2 suggests an interpretation of the FDI in terms of price misperception, rather than departure from first-order conditions. In this view, the consumer properly optimizes, but using prices in the vicinity of the true ones. Though narrower in its interpretation, it has the advantage of remaining meaningful without requiring quasi-concavity. Let the *Price Misperception Index* (PMI) of $\mathcal{D}$ be the infimum of all ε's such that (3a) and (3b) holds for some continuous and strictly monotone utility function v and some perceived price vectors $(p^c(p,x))_{(p,x)\in\mathcal{D}}$ used by the consumer. Quasi-concavity of v is *not* required. Nonetheless, Proposition 2 tells us that $\text{PMI}(\mathcal{D}) = \text{FDI}(\mathcal{D})$.

Of course, the PMI can be adapted to any subclass $\mathcal{U}$ of utility functions that are continuous and strictly monotone, simply by requiring v to belong to $\mathcal{U}$. By definition, $\text{PMI}_{\mathcal{U}}(\mathcal{D}) = \varepsilon$ if, and only if, for all $\varepsilon' > \varepsilon$, one can find price vectors $(p^c(p,x))_{(p,x)\in\mathcal{D}}$ such that (7) holds and the modified demand data $\mathcal{D}^c = \{(p^c(p,x),x)|(p,x) \in \mathcal{D}\}$ is rationalizable with respect to $\mathcal{U}$. This has two important implications. First, $\text{PMI}_{\mathcal{U}}(\mathcal{D}) = \text{FDI}_{\mathcal{U}}(\mathcal{D})$ for any class $\mathcal{U}$ of regular utility functions. Second, the methodology suggested in Proposition 9 applies to the PMI as well. Tests in the spirit of Afriat's inequalities have been recently proposed to check rationalizability with respect to important classes of preferences that need not be convex (e.g., expected utility without risk aversion among others, see Polisson, Quah and Renou (2019)). These tests extend at once to accommodate price misperception (by adding consumer price variables), which allows one to determine the PMI.

We conclude by considering the PMI in our application from Section 5.2. It turns out that our theoretical results, combined with our earlier empirical ones, provide a very similar picture of the data. We already know that the PMI of the dataset, which does not require quasi-concavity, is identical to the FDI of the dataset. Next consider state independence. The astute reader would point out that state independence implies risk aversion with convex preferences. Indeed, getting $(x_1 + x_2)/2$ in both states must be preferred to the bundle (x_1, x_2). Could it be that risk aversion plays a critical role in the data analysis? Notice first that there is no evidence of risk-loving behavior in the data.[34] That being said, one might entertain the possibility of intermediate risk preferences. For instance, considering a parametric class of Gul (1991) preferences capturing disappointment aversion, Halevy et al. (2018) finds in another dataset that some subjects are best described via a negative parameter of disappointment aversion (in which case they are elation loving). Such preferences are neither convex nor concave. Thus we may be interested in the PMI of the class $\mathcal{SI}^-$, obtained by *dropping quasi-concavity* from $\mathcal{SI}$.

It turns out that $\mathrm{PMI}_{\mathcal{SI}^-}$ is in fact equal to $\mathrm{FDI}_{\mathcal{SI}}$. Indeed, we have:

$$\mathrm{FDI}_{\mathcal{SI}}(\mathcal{D}) = \mathrm{FDI}(\bar{\mathcal{D}}) = \mathrm{PMI}(\bar{\mathcal{D}}) = \mathrm{PMI}_{\mathcal{SI}^-}(\mathcal{D}), \tag{9}$$

where $\bar{\mathcal{D}}$ is the mirror-extended dataset from (8). The first two equalities are by now familiar. The first was already applied in Section 5.2, and is established in the Appendix. The second is a corollary of Proposition 2. As for the third equality, observe that if demand data is rationalizable by a utility function in $\mathcal{SI}^-$, then its mirror-extended version is rationalizable by a continuous and strictly increasing utility function. Applying this fact to the modified demand data in the definition of the PMI, it follows that $\mathrm{PMI}(\bar{\mathcal{D}}) \leq \mathrm{PMI}_{\mathcal{SI}^-}(\mathcal{D})$. The reverse inequality also holds, since if a mirror-extended dataset is rationalizable by a continuous and strictly increasing u, then the original data is rationalizable by a utility function $v \in \mathcal{SI}^-$, where $v(x) = u(x)$ if $x_2 \leq x_1$ and $v(x) = u(x_2, x_1)$ if $x_2 \geq x_1$.

While dropping quasi-concavity makes no difference for state-independent utility functions, it does make a difference for expected utility (one can use Example 2 of Polisson et al. (2019) to show this). Letting $\mathcal{EU}$ comprise the class of continuous and

[34] A risk-loving consumer would pick a corner solution in each budget set she faces. More than 80% of the subjects never pick a corner solution, more than 95% of the subjects pick a corner solution in less than half the budget sets they face, and only 6 subjects (out of 1182) picked a corner solution in more than 90% of the budget sets they faced.

strictly increasing expected utility functions, $\mathrm{PMI}_{\mathcal{EU}}$ can easily be computed by linear programming through a simple adjustment of Polisson et al. (2019)'s test, both for the true data and the randomly-generated Bronars choices. In this case, however, we can conclude even without computing it that the empirical CDF of $\mathrm{PMI}_{\mathcal{EU}}$ will have to squeeze in the tight space between the red and green curves in Figure 3. To see this, observe that:

$$\mathrm{PMI}_{\mathcal{SI}^-}(\mathcal{D}) \leq \mathrm{PMI}_{\mathcal{EU}}(\mathcal{D}) \leq \mathrm{PMI}_{\mathcal{EU}^r}(\mathcal{D}) = \mathrm{FDI}_{\mathcal{EU}^r}(\mathcal{D}). \tag{10}$$

The first two inequalities follows from the fact that $\mathcal{EU}^r \subset \mathcal{EU} \subset \mathcal{SI}^-$, while the equality follows from the fact that the FDI and PMI must coincide for any class of regular utility functions. Thus, combining (9) and (10), we see that $\mathrm{FDI}_{\mathcal{SI}} = \mathrm{PMI}_{\mathcal{SI}^-} \leq \mathrm{PMI}_{\mathcal{EU}} \leq \mathrm{FDI}_{\mathcal{EU}^r}$, which reinforces our conclusion from Section 5.2 using the PMI and without relying on quasi-concavity.

References

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

Afriat, S.N. (1973), On a System of Inequalities in Demand Analysis: An Extension of the Classical Method, *International Economic Review*, 14(2), 460–472.

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

Becker, G. (1962), Irrational Behavior and Economic Theory, *Journal of Political Economy*, 70, 1–13.

Bronars, S. (1987), The Power of Nonparametric Tests of Preference Maximization, *Econometrica*, 55(3), 693–698.

Brown, D. and C. Calsamiglia (2007), The Nonparametric Approach to Applied Welfare Analysis, *Economic Theory*, 37, 183–188.

Chiappori P.A. and J.C. Rochet (1987), Revealed Preferences and Differentiable Demand, *Econometrica*, 55(3), 687–691.

Choi, S., S. Kariv, W. Müller, and D. Silverman (2014), Who Is (More) Rational?, *American Economic Review*, 104, 1518–1550.

de Clippel, G., and K. Rozen (2018), Bounded Rationality and Limited Datasets, *mimeo.*

Debreu, G. and T.C. Koopmans (1982), Additively Decomposed Quasiconvex Functions, *Mathematical Programming*, 24(1), 1–38.

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

Diewert, E. (2012), Afriat's Theorem and Some Extensions to Choice under Uncertainty, *Economic Journal*, 122, 305–331.

Echenique, F., T. Imai and K. Saito (2019), Approximate Expected Utility Rationalization, *mimeo.*

Echenique, F., S. Lee and M. Shum (2011), The Money Pump as a Measure of Revealed Preference Violations, *Journal of Political Economy*, 119(6), 1201–1223.

Esponda, I. and D. Pouzo (2016), Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models, *Econometrica*, 84(3), 1093–1130.

Forges, F. and E. Minelli (2009), Afriat's Theorem for General Budget Sets, *Journal of Economic Theory*, 144, 135–145.

Gabaix, X. (2014), A Sparsity-Based Model of Bounded Rationality, *The Quarterly Journal of Economics*, 129(4), 1661–1710.

Green, R. and S. Srivastava (1986), Expected Utility Maximization and Demand Behavior, *Journal of Economic Theory*, 38, 313–323.

Gul, F. (1991), A Theory of Disappointment Aversion, *Econometrica*, 59(3), 667–686.

Halevy, Y., D. Persitz and L. Zrill (2018), Parametric Recoverability of Preferences, *Journal of Political Economy*, 126 (4), 1558–1593.

Ito, K. (2014), Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing, *American Economic Review*, 104(2), 537–563.

Lee, P., and K-C. Wong (2005), Revealed preference and differentiable demand, *Economic Theory*, 25, 855–870.

Matzkin, R. and M. Richter (1991), Testing Strictly Concave Rationality, *Journal of Economic Theory*, 53, 287–303.

POLISSON, M., J. QUAH AND L. RENOU (2019), Revealed Preferences over Risk and Uncertainty, *American Economic Review*, forthcoming.

ROSE, H. (1958), Consistency of Preference: The Two-Commodity Case, *Review of Economic Studies*, 25, 124–125.

RUBINSTEIN, A. AND Y. SALANT (2012), Eliciting Welfare Preferences from Behavioral Data Sets, *Review of Economic Studies*, 79, 375–387.

SAEZ, E., J. SLEMROD AND S. GIERTZ (2012), The Elasticity of Taxable Income with Respect to Marginal Tax Rates:A Critical Review, *Journal of Economic Literature*, 50(1), 3–50.

STEINER, J. AND C. STEWART (2016), Perceiving Prospects Properly, *American Economic Review*, 106, 1601–1631.

VARIAN, H.R. (1982), The Nonparametric Approach to Demand Analysis, *Econometrica*, 50(4), 945–973.

VARIAN, H.R. (1990), Goodness-of-Fit in Optimizing Models, *Journal of Econometrics*, 46, 125–140.

WOODFORD, M. (2012), Prospect Theory as Efficient Perceptual Distortion, *American Economic Review*, 102(3), 41–46.

Appendix

Proof of Proposition 1 (Ordering satisfying the axioms from Section 2.1)

Let $\alpha = 1/x_2$, $\beta = 1/y_2$, $\alpha' = 1/x_2'$ and $\beta' = 1/y_2'$. Then, Unit Invariance tells us that

$$(x, y) \succeq (x', y') \Leftrightarrow ((x_1/x_2, 1), (y_1/y_2, 1)) \succeq ((x_1'/x_2', 1), (y_1'/y_2', 1)).$$

Using Measurement Invariance with $\alpha = y_2/y_1$ and $\alpha' = y_2'/y_1'$, this means

$$(x, y) \succeq (x', y') \Leftrightarrow \left(\left(\frac{x_1/x_2}{y_1/y_2}, 1\right), (1, 1)\right) \succeq \left(\left(\frac{x_1'/x_2'}{y_1'/y_2'}, 1\right), (1, 1)\right).$$

Regularity implies the existence of a function $f : \mathbb{R}^2_{++} \times \mathbb{R}^2_{++} \to \mathbb{R}$ such that $(x, y) \succeq (x', y')$ if and only if $f(x, y) \geq f(x', y')$. Defining $g(\alpha) := f((\alpha, 1), (1, 1))$, we have $(x, y) \succeq (x', y')$ if and only if $g(\frac{x_1/x_2}{y_1/y_2}) \geq g(\frac{x_1'/x_2'}{y_1'/y_2'})$. By Monotonicity, g strictly increases for $\alpha \geq 1$. By Symmetry, it must be that $g(\alpha) = g(1/\alpha)$ and thus $g(\alpha) = g(\max\{\alpha, 1/\alpha\})$. As $\max\{\alpha, 1/\alpha\} \geq 1$ for all α, the result follows from inverting g. *Q.E.D.*

We now establish a few lemmas that we will use to prove Proposition 2.

Lemma 1 *Let $A \subset \mathbb{R}^L_{++}$ be a finite set of bundles. The following statements hold:*

(a) There exists a price vector $q \in \mathbb{R}^L_{++}$ and a bundle $x \in A$ such that $q \cdot x > q \cdot y$ for all $y \in A \setminus \{x\}$;

(b) There exists a price vector $q \in \mathbb{R}^L_{++}$ such that $q \cdot x \neq q \cdot y$ for all $x, y \in A$.

Proof. (a) Let x be the unique maximal element of A according to the following lexicographic order: x has the largest quantity of good 1 among all bundles in A, it has the largest quantity of good 2 among all the bundles singled out so far, etc. Let $q = (1, \eta, \eta^2, \ldots, \eta^{L-1})$. We claim that *(a)* is satisfied using x and the price vector q, provided that $\eta > 0$ is small enough. Indeed, for each alternative $y \in A$, let ℓ be the first good for which $y_\ell \neq x_\ell$. Then, by definition of x, $y_\ell < x_\ell$. Then $q \cdot (x - y) = \eta^{\ell-1}[(x_\ell - y_\ell) + \sum_{k=\ell+1}^{L} \eta^{k-\ell}(x_k - y_k)]$. Clearly, there exists a threshold η^* such that this expression is strictly positive for all $\eta < \eta^*$, which proves *(a)*.

(b) Pick x and q as in *(a)*. Call them x^1 and q^1. If $q^1 \cdot y \neq q^1 \cdot z$, for all pairs y, z of distinct bundles in $A \setminus \{x^1\}$, then we are done. Otherwise, we show that there exists another price vector q^2 that preserves expenditure comparisons under q^1 when they

are strict, but breaks some ties: (i) $q^1 \cdot y > q^1 \cdot z \Rightarrow q^2 \cdot y > q^2 \cdot z$ for all $y, z \in A$, and (ii) there exists $y, z \in A$ such that $q^1 \cdot y = q^1 \cdot z$ and $q^2 \cdot y \neq q^2 \cdot z$. Property (b) will then follow at once by iterating this reasoning. Let $B \subset A$ be a set of at least two elements such that such that $q^1 \cdot y = q^1 \cdot z$ for all $y, z \in B$. By applying (a), there exists $x^2 \in B$ and p such that $p \cdot x^2 > p \cdot y$ for all $y \in B \setminus \{x^2\}$. Let then $q^2 = q^1 + \eta p$. We check that (i) and (ii) are satisfied by picking some $\eta > 0$ small enough. For any $y, z \in A$, $q^2 \cdot (y - z) = q^1 \cdot (y - z) + \eta p \cdot (y - z)$. So there is a threshold η^* such that (i) holds for all $\eta < \eta^*$. By construction, $q^2 \cdot x^2 > q^2 \cdot y$, for all $y \in B$ other than x^2 and so (ii) is satisfied too, as desired. *Q.E.D.*

Lemma 2 *If $\mathcal{D}$ satisfies GARP, then for each $(p, x) \in \mathcal{D}$ there exists a sequence of price vectors $(p^c_n(p, x))_{n \geq 1}$ in $\mathbb{R}^L_{++}$ such that (a) $p^c_n(p, x) \to p$ as $n \to \infty$, and (b) $\mathcal{D}_n = \{(p^c_n(p, x), x) | (p, x) \in \mathcal{D}\}$ satisfies SARP for all $n \geq 1$.*

Proof. Let $q \in \mathbb{R}^L_{++}$ be a fictitious price vector such that $q \cdot x \neq q \cdot y$ for all bundles $x, y \in X = \{x \in \mathbb{R}^L_+ \mid (p, x) \in \mathcal{D} \text{ for some } p \in \mathbb{R}^L_{++}\}$ (see (b) in Lemma 1). For each $(p, x) \in \mathcal{D}$, consider a new price vector $p^c_\eta(p, x) = \eta q + (1 - \eta) p$, where $\eta > 0$ is small enough that $p^c_\eta(p, x) \cdot x' > p^c_\eta(p, x) \cdot x$, for all (p, x) and (p', x') in $\mathcal{D}$ with $p \cdot x' > p \cdot x$. We now check that $\mathcal{D}_\eta = \{(p^c_\eta(p, x), x) | (p, x) \in \mathcal{D}\}$ satisfies SARP. Suppose there is a sequence $\{(p^k, x^k) | 1 \leq k \leq K\}$ in $\mathcal{D}$ such that $x^{k+1} \neq x^k$ and $p^c_\eta(p^k, x^k) \cdot x^k \geq p^c_\eta(p^k, x^k) \cdot x^{k+1}$ for all $1 \leq k \leq K$ (with the convention $K + 1 = 1$). First, observe $p^k \cdot x^k \geq p^k \cdot x^{k+1}$ for all $1 \leq k \leq K$, by definition of p^c_η. Now, suppose that $p^k \cdot x^k = p^k \cdot x^{k+1}$ for all $1 \leq k \leq K$. Then $q \cdot x^k \geq q \cdot x^{k+1}$, for all $1 \leq k \leq K$. Furthermore, the inequalities must be strict, by definition of q. We get $q \cdot x^1 > q \cdot x^2 > \cdots > q \cdot x^K > q \cdot x^1$, a contradiction. Thus the existence of such a sequence $\{(p^k, x^k) | 1 \leq k \leq K\}$ means that $\mathcal{D}'$ violates GARP, a contradiction. It must be that $\mathcal{D}_\eta$ satisfies SARP, which concludes the proof of this lemma. *Q.E.D.*

Say that $\mathcal{D}$ with ε-rationalizable with strict inequalities if Definition 1 holds with strict inequalities in (2) from the paper, instead of weak.

Lemma 3 *ε-rationalizability is equivalent to ε-rationalizability with strict inequalities.*

Proof. Sufficiency is obvious. For necessity, suppose that the regular utility function u ε-rationalizes $\mathcal{D}$. Definition 1 specifies a vector $g(p, x) \in \partial u(x)$ for each $(p, x) \in \mathcal{D}$. Hence x is u-maximal in the fictitious budget set arising when taking $g(p, x)$ as a price vector. The fictitious dataset $\{(g(p,x),x) \mid (p,x) \in \mathcal{D}\}$ then satisfies SARP. The modified dataset $\mathcal{D}' = \{(\alpha p + (1-\alpha)g(p,x) \mid (p,x) \in \mathcal{D}\}$ also satisfies SARP for all $\alpha > 0$ small enough. Indeed, should α be small enough, $(\alpha p+(1-\alpha)g(p,x))\cdot(x'-x) > 0$ for all $(p,x) \in \mathcal{D}$ and $x' \in X \setminus \{x\}$ such that $g(p,x)\cdot(x'-x) > 0$ (that is, the modified dataset adds no new revealed preference comparison). There is thus a regular utility function v that rationalizes $\mathcal{D}'$, which means that

$$\hat{g}(p,x) = \alpha p + (1-\alpha)g(p,x) \in \partial v(x),$$

for all $(p,x) \in \mathcal{D}$. Hence, for any pair ℓ, ℓ' of goods, we have:

$$(1-\varepsilon)\frac{p_\ell}{p_{\ell'}}\hat{g}_{\ell'}(p,x) = (1-\varepsilon)[\alpha p_\ell + (1-\alpha)g_{\ell'}(p,x)\frac{p_\ell}{p_{\ell'}}]$$

$$\leq (1-\varepsilon)\alpha p_\ell + (1-\alpha)g_\ell(p,x) = \hat{g}_\ell(p,x) - \alpha\varepsilon p_\ell < \hat{g}_\ell(p,x),$$

where the weak inequality follows from the definition of g. This proves that v ε-rationalizes $\mathcal{D}$ with strict inequalities, as desired. *Q.E.D.*

Proof of Proposition 2 (ε-Rationalizability and misperceived prices)

(Necessity) See discussion in the paragraph after the statement of the proposition, which relied on Lemma 3 proved above.

(Sufficiency) For each $(p,x) \in \mathcal{D}$, let $p^c(p,x)$ be a price vector satisfying (3b) such that (3a) holds. Since v rationalizes the demand data $\mathcal{D}' = \{(p^c(p,x),x)|(p,x) \in \mathcal{D}\}$, we know $\mathcal{D}'$ satisfies GARP. Using Lemma 2, we can adjust price vectors appearing in $\mathcal{D}'$ a little bit to satisfy SARP, while still satisfying (3b). By Lee and Wong (2005), there exists a continuous, strictly concave and strictly monotone utility function rationalizing this modified demand data, which means that $\mathcal{D}$ is ε-rationalizable. *Q.E.D.*

Lemma 4 (Linearity of the superdifferential) *Take $f_1, \ldots, f_L : \mathbb{R}^L_{++} \to (-\infty, \infty)$ strictly concave and $(\beta_1, \ldots, \beta_L) \in \mathbb{R}^L_{++}$. Then*

(i) $\partial(\sum_{\ell=1}^L \beta_\ell f_\ell) = \sum_{\ell=1}^L \beta_\ell \partial f_\ell$, where the RHS is given by pointwise addition.

(ii) Suppose that for all ℓ there is $\hat{f}_\ell : \mathbb{R}_+ \to \mathbb{R}$ such that $f_\ell(x) = \hat{f}_\ell(x_\ell)$ for all x. Then $g \in \partial(\sum_{\ell=1}^L f_\ell)(x)$ if and only if $(\beta_1 g_1, \ldots, \beta_L g_L) \in \partial(\sum_{\ell=1}^L \beta_\ell f_\ell)(x)$.

Proof. (i) A vector g is a supergradient of f_ℓ at x if $f_\ell(x) \geq f_\ell(x') + (x' - x) \cdot g$ for all x'. Multiplying by $\beta_\ell > 0$, a first observation is that g is a supergradient of f_ℓ at x if and only $\beta_\ell g$ is a supergradient of $\beta_\ell f_\ell$ at x; that is, $\partial(\beta_\ell f_\ell)(x) = \beta_\ell \partial f_\ell(x)$. Moreover, since g is a subgradient of a convex function h at x if $h(x') \geq h(x) + (x' - x) \cdot g$ for all x', a second observation is that g is a supergradient of $-h$ if and only if it is a subgradient of h. Following Rockafeller (1970), each $-\beta_\ell f_\ell$ may be extended to a proper convex function on $\mathbb{R}^L$. As the effective domains of the functions coincide, $\partial(-\sum_{\ell=1}^L \beta_\ell f_\ell) = \sum_{\ell=1}^L \partial(-\beta_\ell f_\ell)$ by Rockafeller's Theorem 23.8. The first two observations then complete the proof.
(ii) The desired result follows from (i) and the following claim, which we will prove: under the conditions stated in (ii), we have $g \in \partial(\sum_{\ell=1}^L \beta_\ell f_\ell)(x)$ if and only if for every ℓ, the vector which has ℓ-th component equal to g_ℓ and is zero otherwise, belongs to $\beta_\ell \partial f_\ell(x)$. To prove this claim, first observe that by (i) above, $g \in \partial(\sum_{\ell=1}^L \beta_\ell f_\ell)(x)$ if and only if $g = \sum_{\ell=1}^L \hat{g}^\ell$, where $\hat{g}^\ell \in \partial(\beta_\ell f_\ell)(x)$. Since $\hat{g}^\ell$ by definition satisfies $\beta_\ell f_\ell(x) \geq \beta_\ell f_\ell(x') + (x' - x) \cdot \hat{g}^\ell$ for all x', including for all x' with $x_\ell = x'_\ell$ but $x \neq x'$, it must be that $\hat{g}^\ell_{\ell'} = 0$ for all $\ell' \neq \ell$. *Q.E.D.*

Proof of Proposition 3 (ε-Rationalizability and additively separable utility)

We begin with some preliminary definitions. For each ℓ, let $\bar{u}_\ell : \mathbb{R}^L_+ \to \mathbb{R}$ be defined by $\bar{u}_\ell(x) = u_\ell(x_\ell)$. Given any $\beta : \mathcal{D} \to \mathbb{R}_{++}$, let $\bar{U}(y|x, \beta) = \sum_{\ell=1}^L \beta_\ell(p, x)\bar{u}_\ell(y)$ be the perturbed additive utility function. Let $\bar{U}(y) = \sum_{\ell=1}^L \bar{u}_\ell(y)$.

We now show necessity. Given regularity, note that condition (4a) holds for $x \in \mathbb{R}^L_+$ and some $\beta : \mathcal{D} \to \mathbb{R}_{++}$ if and only if the first-order optimality condition $p \in \partial\bar{U}(x|x, \beta)$ holds for all $(p, x) \in \mathcal{D}$. By Lemma 4, this holds if and only if for some $g \in \partial\bar{U}(x)$, $p_\ell = \beta_\ell g_\ell$ for all ℓ. Thus optimality for the additive utility function in (4a) holds by setting $\beta_\ell(p, x) = p_\ell / g_\ell$, where $g \in \partial\bar{U}(x)$ is chosen to satisfy the defining bounds of ε-Rationalizability using $\bar{U}(x)$ at (p, x). Then (4b) follows.

For sufficiency, (4a) implies that for each $(p, x) \in \mathcal{D}$, the optimality condition $p \in \partial\bar{U}(x|x, \beta)$ holds for all $(p, x) \in \mathcal{D}$. By Lemma 4, this holds if and only if for some $g \in \partial\bar{U}(x)$, $p_\ell = \beta_\ell g_\ell$ for all ℓ. The desired result then follows from (4b). *Q.E.D.*

Proof of Proposition 4 (ε-rationalizability and acyclic satisfiability)

It remains to prove sufficiency. By acyclic satisfiability, there is a strict ordering $\succ$ over X with the feature that for all $(p, x) \in \mathcal{D}$, there is $g(p, x) \in C_\varepsilon(p, x)$ such that $x \succ x'$ for all $x' \in \Gamma(x, g(p, x))$. We now construct an auxiliary demand data

$\mathcal{D}' = (g(p,x),x)_{(p,x)\in\mathcal{D}}$ where the vector $g(p,x)$, which is strictly positive, is taken to be the price vector when x is chosen. This data satisfies SARP, since any cycle in the revealed preferences from $\mathcal{D}'$ would imply a cycle in $\succ$, a contradiction. Thus by Lee and Wong (2005) there exists a continuous, strictly monotone and strictly quasi-concave (i.e., regular) function $u : \mathbb{R}^L_+ \to \mathbb{R}$ that rationalizes $\mathcal{D}'$ in the classic sense. In particular, for any $(g(p,x),x) \in \mathcal{D}'$, optimality of the demands requires $g(p,x) \in \partial u(x)$. Since $g(p,x) \in C_\varepsilon(p,x)$, we know by construction that for all $(p,x) \in \mathcal{D}$,

$$\frac{p_\ell}{p_{\ell'}}(1-\epsilon) \leq \frac{g_\ell(p,x)}{g_{\ell'}(p,x)} \leq \frac{p_\ell}{p_{\ell'}}\left(\frac{1}{1-\epsilon}\right).$$

Hence the original demand data $\mathcal{D}$ is ε-rationalizable. *Q.E.D.*

Proof of Proposition 5 (FDI with Two Commodities and Two Observations)

Proof. Clearly, the FDI is zero in the absence of SARP violation. Suppose instead that there is a SARP violation (implying $x \neq x'$), and that $\delta(p,o(x,x')) \leq \delta(p',o(x,x'))$ (a similar argument applies if the opposite inequality holds). Following the reasoning from Section 3, it is easy to check that the demand data is ε-rationalizable if and only if there exists $(p^y,y) \in \{(p,x),(p',x')\}$ and $g \in C_\varepsilon(p^y,y)$ such that $g\cdot y' > g\cdot y$, where y' is the element in $\{x,x'\}$ distinct from y. (In words, one of the two budget lines can be adjusted within the limits imposed by ε in a way that eliminates the SARP violation.) First, we check that the demand data is ε-rationalizable for all $\varepsilon > f^{-1}(\delta(p,o(x,x')))$, where $f(x) := 1/(1-x)$ for all x. Notice that $o(x,x')\cdot x = o(x,x')\cdot x'$, by construction, and hence one can find a vector $g(p,x)$ such that $g(p,x)\cdot x' > g(p,x)\cdot x$ and $g(p,x)$ close enough to $o(x,x')$ that $\varepsilon \geq f^{-1}(\delta(p,g(p,x)))$, or $\frac{1}{1-\varepsilon} \geq \delta(p,g(p,x))$. Second, consider an $\varepsilon \leq f^{-1}(\delta(p,o(x,x')))$. It is not difficult then to check that $g\cdot x' \leq g\cdot x$ and $g'\cdot x \leq v'\cdot x'$, for all $g,g' \in \mathbb{R}^2_{++}$ such that $\delta(g,p) \leq \frac{1}{1-\varepsilon}$ and $\delta(g',p') \leq \frac{1}{1-\varepsilon}$. Hence $\mathcal{D}$ is not ε-rationalizable. The above arguments show that the FDI is $f^{-1}(\delta(p,o(x,x')))$, from which the result follows. *Q.E.D.*

Proof of Proposition 6 (FDI with Two Commodities)

Proof. Clearly, if a demand data is ε-rationalizable, then so is any pair of observations in it. Hence, the FDI of any demand data if larger or equal than the FDI associated to any pair of observations. As for the converse, we show that the demand data is ε-rationalizable, for any ε strictly larger than the FDI's associated to all pairs of observations.

By Proposition 4, it is sufficient to prove that $\mathcal{R}_\varepsilon(\mathcal{D})$ is acyclically satisfiable. We do that by induction: for each subset X' of remaining elements in X, there is an element $x \in X'$ and $v \in C_\varepsilon(p, x)$ such that $\Gamma(x', v) \cap X' = \emptyset$.

Fix the set $X' \subseteq X$ of remaining elements in X. Let X'' be the subset of elements y in X' such that y cannot be obtained as a convex combination of elements in $X' \setminus \{y\}$. If X'' is a singleton, say $X'' = \{y\}$, then $z \geq y$ for all $z \in X'$. Then $\Gamma(y, v) \cap X' = \emptyset$, for all $v \in \mathbb{R}^2_{++}$, and in particular for any $v \in C_\varepsilon(p, y)$, whatever p such that $(p, y) \in \mathcal{D}$. Next, suppose that X'' contains two or more elements. Let's enumerate X'': $X'' = \{x^1, x^2, \ldots, x^K\}$ with the property that $x^j_1 < x^k_1$ for all $j < k$ (in which case we also have $x^j_2 > x^k_2$). If there is $x \in X''$ such that $p \cdot x < p \cdot y$ for all $y \in X'$ and all p such that $(p, x) \in \mathcal{D}$, then $\Gamma(x, p) \cap X' = \emptyset$, and we are done since $p \in C_\varepsilon(p, x)$. So let's assume instead that for all k there exists p and $j \neq k$ such that $(p, x^k) \in \mathcal{D}$ and $p \cdot x^j \leq p \cdot x^k$.

By the structure of X'', given any such p, either all j's with this property are larger than k, or all such j's are smaller than k. Of course, when varying the p's such that $(p, x^k) \in \mathcal{D}$, it may be possible to find larger-than-k j's associated to some p's, and smaller-than-k j's associated to others. For each k, let $f(k)$ be the letter L (resp. R) if we find only smaller-than-k (resp. larger-than-k) j's when considering the various p's such that $(p, x) \in \mathcal{D}$. Thus f indicates the direction in which Samuelson revealed preferences can occur. Finally, $f(k) = LR$ if there exists both p and p' such that $(p, x^k) \in \mathcal{D}$, $(p', x^k) \in \mathcal{D}$, $p \cdot x^j \leq p \cdot x^k$ for $j < k$, and $p' \cdot x^j \leq p' \cdot x^k$ for $k < j$.

Suppose there exists j such that $f(x^j) = R$ and $f(x^{j+1}) = L$. Let p^j be the price vector with the smallest price ratio p^j_1/p^j_2 among those such that $(p^j, x^j) \in \mathcal{D}$. Let p^{j+1} be the price vector with the largest price ratio p^{j+1}_1/p^{j+1}_2 among those such that $(p^{j+1}, x^{j+1}) \in \mathcal{D}$. Since the FDI of $\{(p^j, x^j), (p^{j+1}, x^{j+1})\}$ is stricly less than ε, it must be that $\delta(o(x^j, x^{j+1}), p^j) < \frac{1}{1-\varepsilon}$ or $\delta(o(x^j, x^{j+1}), p^{j+1}) < \frac{1}{1-\varepsilon}$. Suppose the latter holds (a similar reasoning holds in the other case), then consider a vector v that is very close to $o(x^j, x^{j+1})$ but with a slightly smaller first component (making the line orthogonal to it slightly flatter). By construction, $\Gamma(x^{j+1}, v) = \emptyset$ and $v \in C_\varepsilon(p, x^{j+1})$, as $\delta(p^{j+1}, v)$ will remain strictly smaller than $\frac{1}{1-\varepsilon}$ when v is close enough to $o(x^j, x^{j+1})$. That same vector v also belongs to $C_\varepsilon(p', x^{j+1})$ for all p' such that $(p', x) \in \mathcal{D}$ and $p'_1/p'_2 \geq o_1(x^j, x^{j+1})/o_2(x^j, x^{j+1})$ (the line orthogonal to p' is steeper than the line orthogonal to $o(x, x')$). Finally, for vectors p' such that $(p', x) \in \mathcal{D}$ and $p'_1/p'_2 < o_1(x^j, x^{j+1})/o_2(x^j, x^{j+1})$, $\Gamma(x^{j+1}, p') = \emptyset$ (and, trivially, $p' \in C_\varepsilon(p', x^{j+1})$)

since $f(x^{j+1}) = L$.

Now notice that $f(x^1) = R$ and $f(x^K) = L$. If there is no j such that $f(x^j) = R$ and $f(x^{j+1}) = L$, then there must be a sequence of indices $\{j, j+1, \ldots, k\}$ such that $f(j) = R$, $f(k) = L$ and $f(i) = LR$ for all $j < i < k$. The reasoning from the previous paragraph can be applied to j and $j+1$. The only issue that may arise is that the price vector q with the smallest price ratio q_1/q_2 among those such that $(q, x^{j+1}) \in \mathcal{D}$ does not admit a $v \in C_\varepsilon(q, x^{j+1})$ such that $\Gamma(x^{j+1}, v) = \emptyset$. (This wasn't an issue previvously, when $f(j+1) = L$, because we could just pick q itself and be sure that $\Gamma(x^{j+1}, q) = \emptyset$. Should $f(j+1) = LR$ instead, $\Gamma(x^{j+1}, q) \neq \emptyset$ and we cannot be sure that ε provides enough leeway to twist the line orthogonal to q in a way that makes the Γ set empty.) Let q'_1/q'_2 be the price ratio in between q_1/q_2 and p_1^{j+1}/p_2^{j+1} such that $\delta(q, q') = \frac{1}{1-\varepsilon}$. If $q' \cdot x^{j+1} < q' \cdot x^i$ for all $i > j+1$, then we are done because $\Gamma(x^{j+1}, v) = \emptyset$ and $v \in C_\varepsilon(q, x^{j+1})$, for some v close enough to q' (obtained by decreasing a little bit the first component).

Assume instead that $q' \cdot x^{j+1} \leq q' \cdot x^i$ for some $i > j+1$ (hence in particular for $i = j+2$). Let p^{j+2} be the price vector with the largest price ratio p_1^{j+2}/p_2^{j+2} among those such that $(p^{j+2}, x^{j+2}) \in \mathcal{D}$. Since the FDI of $\{(q, x^{j+1}), (p^{j+2}, x^{j+2})\}$ is stricly less than ε and $\delta(o(x^{j+1}, x^{j+2}), q) \geq \frac{1}{1-\varepsilon}$, it must be that $\delta(o(x^{j+1}, x^{j+2}), p^{j+2}) < \frac{1}{1-\varepsilon}$. Now the same reasoning as in the above paragraph applies. We would be done if the new q' that arise from it has the property that $q' \cdot x^{j+2} < q' \cdot x^i$ for $i > j+2$. If not, then we iterate the construction. Remember though that $f(x^K) = L$, and thus, even though the q' arising in each step may have failed the property until now, it will have to pass it in the last step. *Q.E.D.*

Proof of Proposition 7 (Relation to CCEI)

Proof. Fix $\varepsilon \geq 0$, and suppose $\text{FDI}(\mathcal{D}) = \varepsilon$. By definition of ε-rationalizability, for every $(p, x) \in \mathcal{D}$, there is $g(p, x) \in C_\varepsilon(p, x)$ such that the auxiliary demand data $\mathcal{D}' = (g(p, x), x)_{(p,x) \in \mathcal{D}}$ is rationalizable (in the classic sense) by a regular utility function u. We will use u to show that Afriat's CCEI for the original data $\mathcal{D}$ is at most $1 - \varepsilon$, which will imply the desired result.

Take any $(p, x) \in \mathcal{D}$, and consider the indifference curve of u passing through x. As illustrated in Figure 6, rationalizability of $\mathcal{D}'$ means the hyperplane determined by the vector $g(p, x)$ (and going through x) separates the upper contour set of x from those bundles below the hyperplane. We claim that $\frac{p \cdot y}{p \cdot x} \geq 1 - \varepsilon$ for any bundle y above

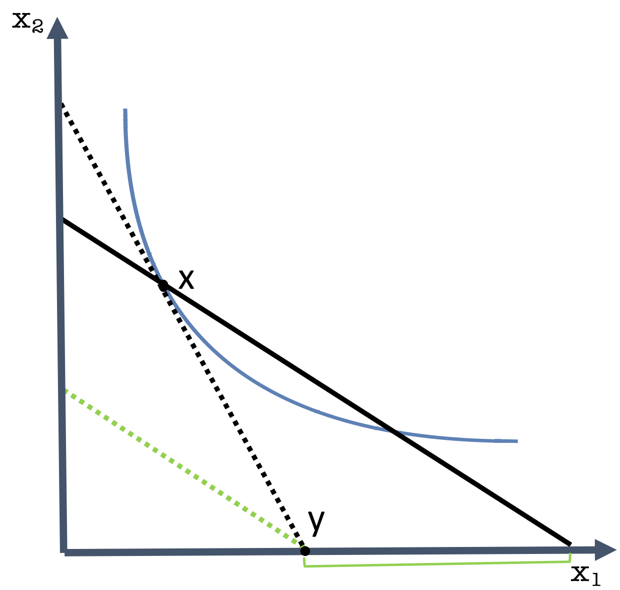

Figure 6: Construction for the proof of Proposition 7.

this hyperplane, i.e., any y such that $g(p,x)\cdot y \geq g(p,x)\cdot x$. This holds trivially if $p = g(p,x)$, so suppose they are different and consider the optimization problem

$$\min_{\{y \mid g(p,x)\cdot y \geq g(p,x)\cdot x\}} p \cdot y.$$

The constraint must bind at the optimum, else the objective could be further reduced. Also, as seen in Figure 6, linearity of the objective and constraint imply the solution must occur at a bundle y with only one positive component: that is, there is ℓ such that $y_\ell = \frac{g(p,x)\cdot x}{g_\ell(p,x)}$ and $y_i = 0$ for all $i \neq \ell$. Using the fact that $g(p,x) \in C_\varepsilon(p,x)$, the minimal expenditure satisfies:

$$p \cdot y = p_\ell y_\ell = p_\ell \sum_{i=1}^{L} \frac{g_i(p,x)}{g_\ell(p,x)} x_i \geq (1-\varepsilon) p_\ell \sum_{i=1}^{L} \frac{p_i}{p_\ell} x_i = (1-\varepsilon) p \cdot x.$$

By quasi-concavity of u, any bundle z with $u(z) \geq u(x)$ must satisfy $g(p,x)\cdot z \geq g(p,x)\cdot x$. Hence, the above inequality shows that if $(1-\varepsilon)$-percent of income is retained, the choice x from the original budget set is strictly preferred under u to all bundles in the remaining budget set. To finish the proof that the CCEI is at most $1-\varepsilon$, observe that any cycle in $\succeq_{A,1-\varepsilon}$ would imply a cycle in the corresponding utilities generated by u, which is impossible. *Q.E.D.*

Proposition 11 (FDI with respect to $\mathcal{SI}$ for $L = 2$) *Let $\mathcal{D}$ be some portfolio demand data with two equally-likely states. Then* $\mathrm{FDI}_{\mathcal{SI}}(\mathcal{D}) = \mathrm{FDI}(\bar{\mathcal{D}})$*, where $\bar{\mathcal{D}}$ is defined in (8).*

Proof. As explained in the text, a utility function is state-independent if, and only if, it is symmetric around the 45^o line, as the two states are equally likely. $\mathrm{FDI}(\overline{\mathcal{D}}) \leq \mathrm{FDI}_{\mathcal{SI}}(\mathcal{D})$ follows from the fact that a symmetric utility function rationalizes (x_2, x_1) given the price vector (p_2, p_1) if it rationalizes x given the price vector p.

For the converse inequality, fix $\varepsilon > \mathrm{FDI}(\overline{\mathcal{D}})$. The proof will be complete after showing that $\mathcal{D}$ is ε-rationalizable with respect to $\mathcal{SI}$. Since $\overline{\mathcal{D}}$ is ε-rationalizable, $\mathcal{R}_\varepsilon(\overline{\mathcal{D}})$ is acyclically satisfiable. There must exist a bundle x among those appearing in $\overline{\mathcal{D}}$ with the following property: for all p such that $(p, x) \in \overline{\mathcal{D}}$, there exists $g \in C_\varepsilon(p, x)$ such that $g \cdot x' > g \cdot x$ for all bundles $x' \neq x$ appearing in $\overline{\mathcal{D}}$. A few observations are worth making: (i) we can assume without loss of generality that $x_2 \leq x_1$, as one can take the symmetric bundle otherwise; (ii) the vector g associated to any (p, x) will have $g_1 < g_2$ if $x_2 < x_1$, as otherwise $g \cdot (x_2, x_1) \leq g \cdot x$; and (iii) the vector g associated to any (p, x) can be chosen so that $v_1 \leq v_2$ if $x_1 = x_2$, since (g_2, g_1) has the same property as g ($\overline{\mathcal{D}}$ is symmetric). Of course, we can eliminate (x_2, x_1) in the next step of the enumeration procedure. We can now iterate. By acyclic satisfiability, there must exist y among the remaining bundles that appear in $\overline{\mathcal{D}}$ with the following property: for all q such that $(q, y) \in \overline{\mathcal{D}}$, there exists $g \in C_\varepsilon(q, y)$ such that $g \cdot y' > g \cdot y$ for all bundles y' appearing in $\overline{\mathcal{D}}$ that are different from x, y, and (x_2, x_1). Similar observations as above apply: we can assume without loss of generality that $y_2 \leq y_1$, the g associated to each (q, y) is such that $g_1 < g_2$ (resp., $g_1 \leq g_2$) if $y_2 < y_1$ (resp., $y_1 = y_2$), and we can eliminate (y_2, y_1) in the next step of the enumeration. Iterating like this, we identify for each $(p, x) \in \mathcal{D}$ with $x_2 \leq x_1$ a strictly positive vector $g(p, x)$ such that $g_1(p, x) \leq g_2(p, x)$ (the inequality being strict if $x_2 < x_1$) and the fictitious demand data

$$\overline{\mathcal{D}}' = \{(g(p, x), x) \mid (p, x) \in \mathcal{D}, x_2 \leq x_1\}$$

satisfies SARP. By Lee and Wong (2005), let u be a continuous, strictly monotone and strictly quasi-concave utility function that rationalizes $\overline{\mathcal{D}}'$.

Let η be such that $1 - \eta > g_1(p, x)/g_2(p, x)$, for all $(p, x) \in \overline{\mathcal{D}}$ such that $x_2 \leq x_1$ and $g_1(p, x) \neq g_2(p, x)$. For each income m, let $k(m)$ be the (unique, by strict quasi-concavity) bundle that is u-maximal in the budget set $\{y \in \mathbb{R}^2_+ \mid y_2 \leq y_1, (1 - \eta)y_1 +$

$y_2 \leq m\}$. As m varies, the function k defines a continuous, monotone path of bundles (an 'income-offer curve'). Let $f : \mathbb{R}_+ \to \mathbb{R}_+$ describe this path: for each x_1, $f(x_1)$ is the unique quantity of good 2 for which one can find m such that $k(m) = (x_1, f(x_1))$. It is easy to check f is well-defined, continuous and strictly increasing.

Let $g : \mathbb{R}_- \to \mathbb{R}_+$ be defined by $g(\alpha) = -\alpha - \eta ln(1-\alpha)$. Notice that $g(0) = 0$, $g'(0) = -1+\eta$, $g'(\alpha) \in (-1, -1+\eta)$ for all $\alpha < 0$, and g is strictly convex. For each bundle x below the diagonal ($x_2 \leq x_1$) and above the income offer curve defined by k ($x_2 \geq f(x_1)$), let $h(x)$ be the unique bundle on that curve such that the translated graph of g, starting at $h(x)$ instead of $(0,0)$, goes through x, that is, such that $x_2 = h_2(x) + g(x_1 - h_1(x))$. To see that this is well-defined, consider the function that associates to any quantity $\beta \geq x_1$ of good 1 the quantity $f(\beta) + g(x_1 - \beta)$. This function equals $f(x_1) \leq x_2$ at $\beta = x_1$ and is strictly increasing. Thus there will be a unique β giving a value x_2.

We are now ready to define a new utility function for bundles below the diagonal:

$$v(x) = \begin{cases} u(x) & \text{if } x_2 \leq f(x_1) \\ u(h(x)) & \text{if } x_2 \geq f(x_1), \end{cases} \tag{11}$$

for each bundle x such that $x_2 \leq x_1$. While it is not difficult to check analytically that v is continuous, strictly monotone and strictly quasi-concave, it is perhaps best to provide a graphical explanation of what the indifference curves and upper contour sets look like for bundles below the diagonal, to understand why these properties hold. Let's focus on an indifference curve going through a bundle x such that $x_2 \leq f(x_1)$ (a similar reasoning holds for the opposite inequality). Hence $v(x) = u(x)$. For bundles y below the income offer curve define by k ($y_2 \leq f(y_1)$), $v(y) = u(y)$, and hence the indifference curve of v coincides with that of u. Let y^* be the bundle at which the u indifference curve going through x crosses the income offer curve. For bundles y above the income offer curve, $v(y) = u(h(y))$. To be indifferent to x according to v, it must thus be that $h(y) = y^*$. Thus the v indifference curve now follows the graph of g translated to start at y^* instead of $(0,0)$. Notice that, by definition of y^*, $(1-\eta, 1)$ is a quasi-gradient of u at y^*. Since the derivative of g at 0 is $-1+\eta$, strict quasi-concavity is preserved. To summarize, indifference curves of v are simply obtained by 'pasting' at points along the income offer curve defined by k indifference curves associated to u below that curve and the translated graph of f above.

By construction, any $g \in \partial v(x)$ at a bundle x with $x_2 < x_1$ has $g_1/g_2 < 1$ (as u has that property if $x_2 < f(x_1)$ and $g'(\alpha) \in (-1, -1+\eta]$ for all $\alpha \leq 0$). Hence the extension of v by symmetry around the 45^o line, $v(x) = v(x_2, x_1)$ for all x such that $x_2 \geq x_1$, remains strictly quasi-concave. Of course, strict monotonicity and continuity are preserved too. Consider now $(p, x) \in \overline{\mathcal{D}}$ such that $x_2 \leq x_1$ and $g_1(p, x) \neq g_2(p, x)$ (which is always the case if $x_2 < x_1$). If y is the unique maximum of v given the fictitious prices $g(p, x)$ and the income $g(p, x) \cdot x$, then $g(p, x)$ belongs to $\partial v(y)$. Then it must be that $y_2 \leq f(x_1)$, as otherwise $\partial v(y)$ contains only rescaling of the gradient of v, which does not match $g(p, x)$ (remember that $1 - \eta > \frac{g_1(p,x)}{g_2(p,x)}$, by definition of η). Either $y_2 < f(y_1)$ and $g(p, x) \in \partial v(y) = \partial u(y)$, or $y_2 = f(y_1)$ and any quasi-gradient $\hat{g} \in \partial v(y)$ such that $1 - \eta \geq \hat{g}_1/\hat{g}_2$, and thus in particular $g(p, x)$, belongs to $\partial u(y)$. Thus y is the u-maximal bundle given the fictitious prices $g(p, x)$ and the income $g(p, x) \cdot x$, that is, $y = x$.

Finally, let $(p, x) \in \overline{\mathcal{D}}$ such that $x_1 = x_2$. If $g_1(p, x) = g_2(p, x)$, then clearly the v-maximal bundle in the budget set with fictitious prices $g(p, x)$ and the income $g(p, x) \cdot x$ is x, by symmetry and strict quasi-concavity of v. If $g_1(p, x) \neq g_2(p, x)$, then the v indifference curve going through x coincides with that of u below the diagonal. Indeed, x is on the income offer curve defined by k, since $1 - \eta > \hat{g}_1/\hat{g}_2$ for *all* quasi-gradient at all bundles $y \neq x$ on the u-indifference curve going through x (indeed, $g(p, x)$ is a subgradient of u at x, $1 - \eta > g_1(p, x)/g_2(p, x)$ by definition of η, and u is strictly quasi-concave). Clearly, $g(p, x)$ thus belongs to $\partial v(x)$, and x is the v-optimal in the budget set with fictitious prices $g(p, x)$ and the income $g(p, x) \cdot x$.

We have thus proved that v rationalizes $\overline{\mathcal{D}}'$, which implies that it ε-rationalizes $\overline{\mathcal{D}}$ with respect to $\mathcal{SI}$, and a fortiori also $\mathcal{D}$. *Q.E.D.*