

RELAXED OPTIMIZATION: HOW CLOSE IS A CONSUMER TO SATISFYING FIRST-ORDER CONDITIONS?*

Geoffroy de Clippel[†] and Kareen Rozen[‡]

Revised July 2020

Abstract

We propose relaxing the first-order conditions in optimization to approximate rational consumer choice. We assess the magnitude of departures with a new, axiomatically-founded measure that admits multiple interpretations. Standard inequality tests of rationality for any given reference class of preferences can be conveniently re-purposed to measure goodness-of-fit with that class. Another advantage of our approach is that it is applicable in any context where the first-order approach is meaningful (e.g., convex budget sets arising from progressive taxation). We apply these ideas to shed new light on existing portfolio-choice data.

*First version July 2018 (under the title ‘Consumer Theory with Misperceived Tastes’). We thank Xavier Gabaix, Andrew Postlewaite, John Quah, attendees at the Workshop on Applications of Revealed Preferences and PennTheon online series, and seminar participants at Brown University, City University of New York, and the University of Pennsylvania 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: karen_rozen@brown.edu.

Contents

1	Introduction	1
2	Consumer Data, ε-Rationalizability and the FDI	4
2.1	Measuring the Discrepancy between a Utility Gradient and a Price Vector	5
2.2	Collections of Price-Vector/Utility-Gradient Pairs	8
2.3	ε -Rationalizability and FOC-Departure Index	9
3	Misperceived Prices	10
4	Misperceived Tastes	13
5	Inequality with Afriat's CCEI	16
6	Computing the PMI and the FDI	18
7	Illustration Using Portfolio Demand Data	20
7.1	Decision-Making Quality	20
7.2	FOSD and Risk-Averse Expected Utility	22
7.3	Plain Expected Utility	25
8	Beyond Linear Budget Sets	26

1 Introduction

In a typical setting, rationality entails making a choice that balances the marginal benefits and costs along all dimensions. Consistency with rationality leaves no room for departure from these first-order conditions (FOCs). One could imagine a more nuanced approach, yielding an approximation of rationality that is graduated by the extent of departure permitted. We explore this idea for consumer choices, where the rational benchmark is captured by the workhorse formula of intermediate microeconomics: *opportunity cost=marginal rate of substitution*. In equivalent mathematical terms, price vectors and utility gradients must be collinear.¹

We propose an axiomatically-founded measure of discrepancy for price vectors and utility gradients, which motivates the following notion of approximate rationality (formalized in Section 2). Given a reference class of utility functions \mathcal{U} and some $\varepsilon \in [0, 1]$, demand data \mathcal{D} is ε -rationalizable given \mathcal{U} if the following inequalities hold for some $u \in \mathcal{U}$:

$$(1) \quad 1 - \varepsilon \leq \frac{MRS_{\ell\ell'}^u(x)}{p_\ell/p_{\ell'}} \leq \frac{1}{1 - \varepsilon},$$

for all observations $(p, x) \in \mathcal{D}$ and all goods $\ell \neq \ell'$. The marginal rates of substitutions (MRS) and opportunity costs must match if $\varepsilon = 0$. The larger ε is, the more permissive ε -Rationalizability becomes. For a given \mathcal{U} , the FOC-Departure Index (FDI) of demand data is the smallest ε (formally an infimum) for which the data is ε -rationalizable given \mathcal{U} . The advantage of considering different reference classes of utility functions is the ability to gauge both how far away the consumer is from simply being rational (that is, applying the definition to the largest reference class of preferences) and the degree of misspecification when considering smaller reference classes.

Our approach defines a new yardstick for measuring the quality of consumer choices while remaining agnostic about the sources of departures from rationality. But the FDI is also meaningful when considering more structured forms of bounded rationality. In

¹More precisely, this formula holds for preferences that are both convex and differentiable, as assumed throughout the introduction for simplicity. Beyond the introduction, we will drop differentiability, which requires slightly more complex notation (for quasi-gradients instead of gradients). Dropping convexity can make the first-order approach inadequate (since the second-order condition is dropped). However, we will see that many of our ideas remain applicable to non-convex preferences, under a related idea of price misperception.

Section 3, we show it measures the minimal degree of *price misperception* required for rationalizing demand data, and offer a new interpretation of the measure as a money-pump multiplier. This interpretation in terms of price misperception also provides a way to extend most results regarding the FDI to preferences that need not be convex. Finally, the FDI can also be understood in terms of *misperception in tastes*. The connection is especially striking when considering additively-separable utility functions, as we show in Section 4. For instance, the FDI given risk-averse expected utility can be interpreted as quantifying misperception of probabilities; and the FDI given exponential discounting can be interpreted as quantifying misperception of discount factors. Furthermore, we establish a tight connection between a central model of reference dependence and ε -rationalizability with respect to additively-separable utility.

Starting with Afriat’s Critical Cost Efficiency Index (CCEI), multiple inconsistency indices have been proposed over the years. Many of them quantify the largest percentage of income that can be retained while restoring rationality (GARP). Afriat’s index, by far the most prevalent, does this while applying the same income-scaling factor to all observations. Varian (1990) and Halevy, Persitz and Zrill (2018) consider alternative aggregators with observation-specific scaling. Echenique, Lee and Shum (2011) aggregates income losses arising from revealed-preference cycles, while Dean and Martin (2015) measures the income loss associated with deleting enough direct revealed-preference comparisons to avoid cycles. Our approach follows an entirely different route left unexplored. Instead of relying on income losses, we measure discrepancies between price vectors and putative utility gradients.²

Yet, perhaps surprisingly, the FDI provides an upper-bound for the percent of lost income according to the CCEI (that is, $1 - \text{CCEI} \leq \text{FDI}$ for all datasets), for any reference class of utility functions. 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. The story is subtler, however, once power is considered: one should take into account whether violations of rationality are likely for the budget sets observed. Using Bronars (1987)’s well-known approach,³ for instance,

²Yet other approaches include Houtman and Maks (1985), that counts how many observations must be dropped to restore consistency, and Apesteguia and Ballester (2015) that assesses how ‘many’ feasible alternatives are superior to chosen options (using the Lebesgue measure in the space of bundles when considering consumer choices).

³The 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.

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. It is thus safer to assess decision-making quality according to both indices.

The FDI presents some advantages. All classic tests of consistency for reference classes of preferences immediately extend to compute the corresponding FDI (see Section 6); the same does not hold for Afriat’s CCEI. The reason for this difference is that demand data is fully rational for *modified prices* under the FDI, while consumers’ choices cease to be optimal in Afriat’s shrunken budget sets (indeed, they are not even feasible). Thus the FDI is a flexible measure, in the sense that it is easily applied for different classes of reference utility functions.

Another advantage of our approach is its portability to problems beyond linear budget sets. The first-order condition characterizes optimality in numerous economic settings. Our approach readily extends to all such settings, by bounding the discrepancy between utility and opportunity-set gradients. Consider, for instance, progressive taxation in labor and investment decisions, which yield convex but non-linear budget sets. Our approach generalizes to this framework by comparing putative utility gradients and *marginal* price vectors. By contrast, it is not always clear how to extend Afriat’s CCEI (or, more generally, money-metric indices) to such settings. There may be multiple, conflicting ways to ‘shrink’ budget sets when prices are non-linear; that is, the general approach of the CCEI becomes ill-defined. Though we mostly focus on textbook consumer theory, we develop this theme further in our concluding section.

We illustrate the applicability of our approach in Section 7, 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 in Section 7.1 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 the FDI *strictly* larger than $1 - \text{CCEI}$ 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 assessment that there is a significant amount of rationality to be found.

In Section 7.2, we shed new light on this data by empirically assessing the validity of different preference restrictions. This portion relates to, but goes beyond, some as-

pects of a contemporaneous and independent paper, Echenique, Imai and Saito (2019). They define a measure of departure from risk-averse expected utility maximization, and apply it to examine the prevalence of such preferences in Choi et al.’s dataset (among others). As it turns out, their measure is a renormalization of FDI_{EUR} , the FOC-Departure Index given the class of risk-averse expected utility. A first advantage of our approach is the ability to compare apples to apples. Echenique et al. empirically contrast their new measure against the CCEI for general preference maximization. We can study the impact of restricting attention to risk-averse expected utility while holding constant the method by which departures are measured. A second advantage of our approach is the opportunity to compute the FDI relative to relevant classes of preferences beyond risk-averse expected utility. Perhaps surprisingly, we find that the empirical departure from risk-averse expected utility maximization is mostly attributable to departure from the plain maximization of a utility function defined over lotteries.⁴ In the spirit of Bronars, we compute the FDI over these different classes of preferences using randomly-generated choices for the budget sets tested in Choi et al. This exercise makes clear that, although their dataset is well-suited for testing basic rationality, richer data (e.g., varying the probability of states) is needed if one hopes to disentangle risk-averse expected utility maximization from the maximization of a quasi-concave utility function defined over lotteries. In Section 7.3, we use the narrower interpretation in terms of price misperception to show these conclusions remain valid even without requiring convex preferences (e.g., plain expected utility instead of risk-averse expected utility).

2 Consumer Data, ε -Rationalizability and the FDI

We observe a consumer selecting consumption bundles 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 convexity (this is relaxed in Sections 3 and 7.3). Such preferences are representable by a *regular* utility function: one that is continuous, strictly monotone, and quasi-concave. The rational benchmark posits that the consumer se-

⁴With equally likely states, treating the bundles as lotteries amounts to state independence or first-order stochastic dominance.

lects bundles through utility maximization over budget sets. If the utility function is differentiable at an interior choice, then price vectors and utility gradients must be collinear (equivalently, opportunity costs must equal marginal rates of substitution at chosen bundles). For expositional convenience, we will focus on this case in the next two subsections, though non-differentiable utility functions and corner solutions will be accommodated when presenting the general definitions afterwards.

2.1 Measuring the Discrepancy between a Utility Gradient and a Price Vector

We propose quantifying departures from rationality by how “close” the consumer is from satisfying the knife-edge collinearity condition. While there are many ways to measure distances between vectors in general (such as the angular or the Euclidean distance), we argue that measuring distances between price vectors and utility gradients should be done a certain way. Indeed, given their economic interpretation, such variables are not uniquely defined. 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. Euclidean distance, however, is sensitive to such rescaling. Moreover, a consumer’s problem is unaffected when modifying how a good’s quantity is measured (e.g., using ounces or grams, gallons or quarts, etc.), provided that prices are adjusted accordingly.⁵ Both angular and Euclidean distance, however, are sensitive to measurement choice.

We start by considering the simpler case of two goods (as in most experimental papers on the topic), and suggest further below a natural extension to accommodate more goods. Let \succeq be a weak ordering over price-vector/utility-gradient pairs:⁶ $(p, g) \succeq (p', g')$ means that “the utility gradient g is farther apart from the price vector p than g' is from p' .” The ultimate goal is to assess degrees of rationality, in which case $(p, g) \succeq (p', g')$ is interpreted as follows: a consumer who picks a bundle at which the utility gradient is g while the price vector is p , is less rational than a consumer

⁵For 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).

⁶While developing our ideas for assessing how far apart a utility gradient is from a price vector, our arguments are equally meaningful for measuring how far apart two price vectors are from each other, and for measuring how far apart two utility gradients are from each other. This will be relevant in Sections 3 and 4 that explore alternative interpretations of the methodology we propose.

who picks a bundle at which the utility gradient is g' while the price vector is p' . Of course, g and g' are not observable (and we will consider possibilities for what they may be), but for the moment we proceed as if we know them. The following axioms capture the invariance properties discussed above.

Unit Invariance $(p, g) \succeq (p', g')$ if, and only if, $(\alpha p, \beta g) \succeq (\alpha' p', \beta' g')$, for all positive scalars $\alpha, \beta, \alpha', \beta'$.

Measurement Invariance $(p, g) \succeq (p', g')$ if, and only if, $((\alpha p_1, p_2), (\alpha g_1, g_2)) \succeq ((\alpha' p'_1, p'_2), (\alpha' g'_1, g'_2))$, for all positive scalars α, α' (and similarly for good 2).

The first axiom reflects the fact that a price vector or a utility gradient is effectively determined only up to a positive linear transformation. It permits different transformations for the different vectors p, p', g and g' , but requires *all dimensions* of the same 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.

Numbering Invariance $(p, g) \sim ((p_2, p_1), (g_2, g_1))$

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

The first axiom ensures existence of a numerical representation. The second means that the distance between a utility gradient and a price vector should be independent of which good is called good 1. The third simply requires that an increase in $\alpha \geq 1$ brings the utility gradient $(\alpha, 1)$ further away from the price vector $(1, 1)$.⁷

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

$$\delta(\hat{p}, \hat{g}) = \max \left\{ \frac{\hat{p}_1/\hat{p}_2}{\hat{g}_1/\hat{g}_2}, \frac{\hat{g}_1/\hat{g}_2}{\hat{p}_1/\hat{p}_2} \right\}$$

for all price vectors \hat{p} and all utility gradients \hat{g} .

⁷Naturally, 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.




Figure 1: Measuring the Discrepancy Between p and g .

Of course, \succeq^* admits many representations: any increasing transformation of δ . Which gets selected is irrelevant for comparing degrees of rationality, as we do later. To fix ideas though, we pick two reference vectors – the price vector $(1, 1)$ and the utility gradient $(1 - \varepsilon, 1)$ – and suggest to assess discrepancies in comparison to how far apart these two reference vectors are from each other. These two vectors are easy to visualize, and provide an intuitive scale where discrepancies are measured by a number between 0 (perfect congruence) and 1 (orthogonal vectors that are in total mismatch). By definition, g is closer to p than $(1 - \varepsilon, 1)$ is from $(1, 1)$ if, and only if,

$$\delta(p, g) = \max\left\{\frac{p_1/p_2}{g_1/g_2}, \frac{g_1/g_2}{p_1/p_2}\right\} \leq \frac{1}{1 - \varepsilon}.$$

The smallest ε satisfying this inequality – $\varepsilon(p, g)$ – is our favored representation of \succeq^* : $\varepsilon(p, g) = 1 - \frac{1}{\delta(p, g)}$. This parametrization happens to be easily interpretable. Consider the scenario depicted on the left panel of Figure 1. The consumer’s indifference curve coincides locally with the segment that is orthogonal to g . While remaining indifferent, the consumer could save money by decreasing his consumption of good 1 by one unit, and increasing his consumption of good 2 by g_1/g_2 units. How many units of good 1 could he then recoup as a ‘freebie’, keeping his original expenditure on x but providing additional utility? Precisely $\varepsilon(p, g)$, since he saved $\$(p_1 - p_2 \frac{g_1}{g_2})$ and good 1 costs $\$p_1$. The two panels on the right of Figure 1 depicts more extreme situations, where $\varepsilon(p, g)$ gets close to 0 or 1.

With more than two goods, we let the discrepancy between p and g be the maximal ε -distance between the projections of p and g on any pair of distinct goods ℓ, ℓ' .

2.2 Collections of Price-Vector/Utility-Gradient Pairs

Our ultimate objective is to measure how close demand data containing *multiple* observations is from being rational. To this end, we must extend \succeq^* , or equivalently our favored representation $\epsilon(\cdot)$, to compare *collections* of price-vector/utility-gradient pairs.⁸ Again, we pretend for the time being that the gradients are known.

Throughout the paper, we will focus on the extension $\epsilon(D) = \max_{(p,g) \in D} \epsilon(p, g)$ for a collection D of price-vector/utility-gradient pairs. The robust lesson from our axiomatic analysis is that gradient ratios over price ratios are the critical objects of study, due to the invariance properties prices and gradients satisfy. These do not imply, however, that one must perform a worst-case analysis when aggregating across observations. Some may argue, for instance, that a consumer who makes highly irrational decisions in rare circumstances is more rational than a consumer with intermediate departures from rationality in a large number of circumstances. A similar aggregation question applies to measures of rationality based on income adjustments needed to eliminate revealed preference cycles.

On that dimension, our max-operator simply follows the worst-case scenario analysis favored by Afriat's (1972, 1973) in his definition of the critical cost efficiency index, which remains so far the predominant index despite alternative aggregators proposed over the years, starting with Varian (1990); see also Halevy, Persitz and Zrill (2018) and references therein. One could also apply those suggested aggregation methods to our $\epsilon(p, g)$: for instance, letting $\epsilon(D)$ be a frequency-weighted average of the measure $\epsilon(p, g)$ for each $(p, g) \in D$. A main reason for the empirical literature's bias towards Afriat's worst-case scenario is the complexity of computation for other aggregation methods, and the same computational issues would arise from applying those other aggregation methods here. Having said this, we note that our notion of ϵ -Rationalizability, and thus the FOC-Departure Index and Price-Misperception Index, can be applied with other aggregators; and it is easy to see from the proofs that our theoretical results either hold verbatim or easily generalize, when using any other monotone aggregation method.⁹

⁸A collection may contain the same pair (p, g) multiple times, as the consumer may pick the same bundle when facing a same budget set multiple times, or be assigned a same gradient at his choices for different budget sets defined by the same price vector p .

⁹That is, with $\epsilon(p, g)$ for singleton collections defined just as above, we can construct other extensions of ϵ to non-singleton collections D besides the max-operator. Say $\epsilon(\cdot)$ is a monotone aggregation method if whenever $D = ((p^i, g^i)_i)$ and $\bar{D} = ((\bar{p}^i, \bar{g}^i)_i)$ satisfy $\epsilon(\bar{p}^i, \bar{g}^i) \geq (\leq) \epsilon(p^i, g^i)$ for all i , then

2.3 ε -Rationalizability and FOC-Departure Index

Now that we understand how to measure departures of price vectors and known utility gradients, we can dispense with the assumption that gradients are known. The notion of ε -rationalizability given a class \mathcal{U} of regular utility functions proceeds as follows. Any given $u \in \mathcal{U}$ defines gradients at demanded bundles, and the question is whether there exists $u \in \mathcal{U}$ that generates a collection D of price-vector/utility-gradient pairs such that $\epsilon(D)$ is lower or equal to the parameter ε .¹⁰

For increased generality, we allow for corner solutions and non-differentiable utility functions (for which normalized gradients are not always uniquely defined). The set of quasi-gradients $\partial u(x)$ is the set of strictly positive vectors defining the supporting hyperplanes of the upper-contour set at x :¹¹

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

At any point x where u is differentiable, $\partial u(x)$ contains a unique vector – the usual gradient $\nabla u(x) = (\frac{\partial u(x)}{\partial x_1}, \dots, \frac{\partial u(x)}{\partial x_L})$ – up to positive rescaling. If \mathcal{U} accommodates non-differentiable utility functions (or if demand data contains corner bundles), then one is free to pick the normalized quasi-gradient compatible with u when checking ε -rationalizability.

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

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

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

For intuition, the ratio $g_\ell / g_{\ell'}$ is uniquely defined for $x \gg 0$ when u is differentiable,

$\epsilon(\bar{D}) \geq (\leq) \epsilon(D)$. Footnotes 10 and 12 describe how the main definitions generalize.

¹⁰Indeed, this description is how one would define ε -Rationalizability when using a more general aggregation method than the max-operator; one must only stipulate that D is constructed such that for all $(p, x) \in \mathcal{D}$ we have $g \in \partial u(x)$, as in Definition 1. The further detail in Definition 1 is specific to the max-operator.

¹¹At the cost of heavier notation later on, one could allow g to belong to $\mathbb{R}_+^L \setminus \{0\}$ 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, to strictly positive price vectors.

and corresponds to the *marginal rate of substitution* (MRS) associated to ℓ and ℓ' :

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

The numerator in (2) simply captures tradeoffs in taste, while accommodating the possibility of multiple implicit utility tradeoffs in case of non-differentiability. The price ratio $p_\ell/p_{\ell'}$ in the denominator represents the opportunity cost. Notice then that (2) simply boils down to the FOC that is characteristic of rationality when $\varepsilon = 0$.

Rationality is an all-or-nothing condition: demand data is either consistent with it or not. In an imperfect world of actual data, it is more useful to have ways to quantify the degree to which data complies with a theory, and ε -Rationalizability naturally lends itself to such measurements.

Definition 2 (FOC-Departure Index) *The FOC-Departure Index of \mathcal{D} given \mathcal{U} , denoted $FDI_{\mathcal{U}}(\mathcal{D})$, is the infimum over all ε such that \mathcal{D} is ε -rationalizable.*

In applications, the reference class \mathcal{U} may include all regular utility functions; but analysts are often interested in special classes of non-parametric preferences, adding requirements such as quasi-linearity, additive separability, homotheticity, expected utility or exponential discounting. Our approach is also applicable for parameter estimations when \mathcal{U} is a parametric class of utility functions (as in the CES family in Fisman, Kariv and Markovits (2007), or the $\beta - \delta$ CRRA class in Balakrishnan, Haushofer and Jakiela (2019), to cite just a couple of examples). The advantage of being able to vary the reference class is apparent from our data analysis in Section 7.

3 Misperceived Prices

Introspection and empirical evidence suggest that people rarely perceive prices perfectly. In the spirit of the Weber-Fechner law, consumers' understanding of prices is likely related to actual prices, but not necessarily a perfect match. They may round up prices to simplify budget arithmetic, and are often subject to systematic biases such as the left-digit effect (Thomas and Morwitz, 2005). People often underestimate the price of add-ons, such as favoring a good with a lower tag price over an alternative that is cheaper once shipping is included (Brown, Hossain and Morgan, 2010). Many buyers also fail to fully incorporate sales taxes in their decisions (Chetty, Looney

and Kroft, 2009). Research in visual perception (Frisby and Stone, 2010) suggests that subjects in experiments who are presented with graphical budget lines can only roughly assess the slopes. Further examples are cited in Gabaix (2014), who proposes a theory of endogenous price misperception based on sparsity-based optimization.

For a consumer who properly optimizes given perceived prices, we would say that the more accurately she perceives prices, the closer she is to being rational. As we pointed out earlier (see Footnote 6), the axiomatic motivation for \succeq^* is equally applicable when comparing price vector pairs: $(p, p^c) \succeq^* (q, q^c)$ means that the perceived price vector p^c is farther apart from the actual price vector p than q^c is from q . A consumer who optimizes using p^c instead of p would be deemed less rational than a consumer who used q^c instead of q . Of course, perceived prices are not directly observable (the same way that utility gradients were not observable in Section 2), but we can figure out the minimal range of ‘price twisting’ required to rationalize a consumer’s choice. For this, fix a class \mathcal{U} (parametric or not) of strictly monotone and continuous utility function.

Definition 3 (Price Misperception Index) *The Price Misperception Index of \mathcal{D} given \mathcal{U} (denoted $PMI_{\mathcal{U}}(\mathcal{D})$) is the infimum over all ε for which one can find $u \in \mathcal{U}$ and associate to all $(p, x) \in \mathcal{D}$ a price vector $p^c \in \mathbb{R}_{++}^L$ such that*

$$(3a) \quad x \in \underset{\{p^c \cdot y \leq p^c \cdot x\}}{\operatorname{argmax}} u(y) \text{ (demanded bundle is } u\text{-maximal under consumer prices),}$$

$$(3b) \quad 1 - \varepsilon \leq \frac{p_\ell/p_{\ell'}}{p_\ell^c/p_{\ell'}^c} \leq \frac{1}{1 - \varepsilon}, \quad \forall \ell \neq \ell' \text{ (consumer prices near true prices).}^{12}$$

As it turns out, the PMI and FDI coincide for any given classes of regular utility functions.

Proposition 2 *Let \mathcal{D} be any demand data, and \mathcal{U} be any class of regular utility functions. Then $PMI_{\mathcal{U}}(\mathcal{D}) = FDI_{\mathcal{U}}(\mathcal{D})$.*

Conceptually, this means price misperception offers a natural, alternative interpretation for the approach developed in the previous section. The result has a simple proof. Demand data is ε -rationalizable given \mathcal{U} if, and only if, there is $u \in \mathcal{U}$ for which ∂u at any demanded bundle has an element that is ε -close to the associated price vector.

¹²For aggregation methods $\epsilon(\cdot)$ besides the max, this condition should be replaced with $\epsilon(D) \leq \varepsilon$, where D is the collection of the true-price vector/consumer-price vector pairs (p, p^c) .

The result follows by interpreting these vectors as the consumer's price vectors.

The PMI has another valuable use. As we all know, the first-order condition provides little information unless accompanied by the second-order condition to guarantee a global optimum. For instance, the FDI with respect to the entire class of strictly monotone and continuous utility functions is zero, as one can always draw indifference curves to match local restrictions at finitely many bundles. Not so for the PMI, which remains meaningful without quasi-concavity. This is useful, for instance, when assessing degrees of consistency with expected utility maximization without imposing risk aversion (see Section 7.3 for an illustration).

Afriat (1967) pointed out that demand data is rationalizable by a regular utility function if, and only if, it is rationalizable by one that need not be quasi-concave. Hence the PMI is invariant whether or not one requires utility functions to be quasi-concave in addition to being continuous and strictly monotone. Combined with Proposition 2, we get a justification for the FDI that does not rely on any quasi-concavity.

Proposition 3 *Let \mathcal{D} be any demand data, \mathcal{U} be the class of all regular utility functions, and \mathcal{U}^- be the larger class obtained by dropping quasi-concavity from the requirements. Then $FDI_{\mathcal{U}}(\mathcal{D}) = PMI_{\mathcal{U}^-}(\mathcal{D})$.*

In addition to the axiomatic justification of δ from Section 2, which also applies when comparing two price vectors (see Footnote 6), Equation (3b) admits another interesting interpretation. A discrepancy between p and p^c means the consumer is susceptible to a money pump. Consider a rational trader with $\$M$ conducting the following trade scheme once. He starts by using his $\$M$ money to buy¹³ 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, assume the consumer accepts trades which leave her budget unchanged.¹⁴ What is the return on investment that the rational trader can make in this scenario?

¹³The consumer's endowment in goods is assumed to be large enough, or M is assumed to be small enough that the consumer can provide the good that the trader wants to buy.

¹⁴If she accepts only trades that strictly increase her budget, then the rational trader 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.

Proposition 4 *The rational trader's return on investment is $\delta(p, p^c)$.*

Proof. The rational trader maximizes $p^c \cdot y$ 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: $p_\ell/p_\ell^c \geq p_k/p_k^c$ 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 $p \cdot y \leq (p_\ell M)/p_\ell^c$. Similar reasoning reveals that an optimal solution is to spend the $\$(p_\ell M)/p_\ell^c$ on a good ℓ' with the highest price ratio when comparing perceived prices to true prices: $p_{\ell'}^c/p_{\ell'} \geq p_k^c/p_k$ for each good k . The profit in that case is $\$(p_{\ell'}^c/p_{\ell'})(p_\ell M)/p_\ell^c$. So

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

is the trader's maximal profit, and $\delta(p, p^c)$ is the return on investment. *Q.E.D.*

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

4 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, and contrary to Section 3, it will be important to restrict attention to convex preferences 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 choice in all circumstances. Alternatively, to save on contemplation and thinking costs, she may merely check her putative choice has no preferable alternative nearby, thereby reaching her choice by tatonnement. Reaching choices through a series of small adjustments may be a reasonable descrip-

tion of the thought process in some circumstances. It is unlikely that every time some prices change, the consumer reassesses her entire budget set, carefully introspects about her preference relation over those bundles, and directly selects the globally preference-maximizing bundle. Conveniently, stepwise and local thinking converges to rational choices when preferences are convex, so long as the consumer correctly assesses her local utility tradeoffs. This is no longer the case if the consumer imperfectly assesses those utility tradeoffs (or if these tradeoffs vary somewhat with factors unobserved by the modeler). Bounding the error in assessing the utility gradient gives rise to ε -rationalizability.

In specific classes of utility functions, our notion of ε -rationalizability can suggest a particular channel for the consumer's misperception (or modeler's misspecification). 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,¹⁵ but also for any additively separable preference. Formally, say u is *additively separable* if $u(x) = \sum_\ell u_\ell(x_\ell)$ for some concave,¹⁶ continuous and strictly monotone utility functions $u_\ell : \mathbb{R}_+ \rightarrow \mathbb{R}$.

Proposition 5 *\mathcal{D} is ε -rationalizable with respect to the additively separable utility function $u(\cdot) = \sum_{\ell=1}^L u_\ell(\cdot)$ if, and only if, for all $(p, x) \in \mathcal{D}$, there is $\beta(p, x) \in \mathbb{R}_{++}^L$ such that*

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

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

¹⁵Technically, the proposition applies to Cobb-Douglas only when restricting attention to strictly positive bundles, as $\log(x_\ell) \in \mathbb{R}$ only if $x_\ell > 0$.

¹⁶Concavity of u_ℓ may seem much stronger than our usual requirement of quasi-concavity for u . 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.

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 once again the δ -function uncovered in Section 2 to measure these discrepancies in preference parameters. In the Cobb-Douglas example, $u_\ell(\cdot) = \alpha_\ell \log(\cdot)$ for each ℓ . For intertemporal choices, with x_ℓ representing consumption at time ℓ , exponential discounting corresponds to the case $u_\ell(\cdot) = \delta^\ell \tilde{u}(\cdot)$ for some time-independent utility function \tilde{u} . Proposition 5 then shows that consumer's errors can be interpreted as misperceived discounting. Similarly, in a setting with risk, where each ℓ is a state of the world, all errors could be attributed to misperceived probabilities.

Reference-dependent choices have become central to behavioral economics since Kahneman and Tversky (1979) and Tversky and Kahneman (1991). We conclude this section by highlighting an interesting relationship between ε -rationalizability and a central model of reference-dependent consumer choices. Say that \mathcal{D} is *consistent with reference-dependent choices for gain-loss parameters* $\gamma \geq 0$ and $\lambda \geq 1$ if one can find continuous, concave and strictly increasing utility functions $u_\ell : \mathbb{R} \rightarrow \mathbb{R}$ and a reference bundle $r(p, x) \in \mathbb{R}_+^L$ for each (p, x) such that, for all $(p, x) \in \mathcal{D}$, the demand x maximizes the modified utility function $v(y|r(p, x))$ over all the bundles $y \in \mathbb{R}_+^L$ such that $p \cdot y \leq p \cdot x$, where

$$v(y|r(p, x)) = \sum_{\ell=1}^L v_\ell(y_\ell|r_\ell(p, x)),$$

$$v_\ell(y_\ell|r_\ell(p, x)) = \begin{cases} u_\ell(y_\ell) + \gamma[u_\ell(y_\ell) - u_\ell(r_\ell(p, x))] & \text{if } y_\ell \geq r_\ell(p, x) \\ u_\ell(y_\ell) + \gamma\lambda[u_\ell(y_\ell) - u_\ell(r_\ell(p, x))] & \text{if } y_\ell \leq r_\ell(p, x) \end{cases}$$

for each $y \in \mathbb{R}_+^L$ and $\ell = 1, \dots, L$. The reader will have recognized the exact same gain-loss utility function (over riskless choices) as in Koszegi and Rabin (2006). The difference is that we remain agnostic as to how the reference point forms: it is viewed as an unobserved parameter that can possibly vary each time the consumer makes a choice. Naturally, having a theory of how reference points arise, or limiting the range in which they may fall would give rise to stronger testable implications.

As it turns out, consistency in this sense is equivalent to ε -rationalizability given an additively-separable utility function, with $\varepsilon = \frac{\gamma(\lambda-1)}{1+\gamma\lambda}$.¹⁷ This means that a consumer

¹⁷Notice how the above model reduces to rationality if, and only if, $\gamma = 0$ or $\lambda = 1$. In that case, ε is indeed equal to zero.

subject to the above form of reference dependence will appear ε -rationalizable for an additively-separable utility function. Vice versa, any dataset that is ε -rationalizable for an additively-separable utility function can be interpreted as data arising from a consumer with gain-loss parameters (γ, λ) . Hence our test of ε -rationalizability given additively separable utility functions (see Section 6 on testing methods) could also be seen as a test for consistency with the above model of reference dependence.

Proposition 6 *\mathcal{D} is consistent with reference-dependent choices for gain-loss parameters $\gamma \geq 0$ and $\lambda \geq 1$ if, and only if, it is ε -rationalizable given an additively separable utility function, with $\varepsilon = \frac{\gamma(\lambda-1)}{1+\gamma\lambda}$.*

5 Inequality with Afriat's CCEI

Afriat's Critical Cost Efficiency Index (CCEI) is the largest percentage of the consumer's budgets that can be retained while satisfying GARP (or, equivalently, eliminating all Samuelson revealed-preference cycles). This notion is immediately applicable when plain rationality is the benchmark model, as characterized by GARP. But when tackling subclasses of utility functions, it is more convenient to follow a different route initiated by Varian (1990) (see also Halevy et al. (2014)). Given any continuous and increasing utility function u , let the *budget efficiency ratio* of an observation $(p, x) \in \mathcal{D}$ be defined as follows:

$$r(p, x; u) = \frac{\min_{\{y \in \mathbb{R}_+^L \mid u(y) \geq u(x)\}} p \cdot y}{p \cdot x}.$$

Then the Critical Cost Efficiency Index of \mathcal{D} given a class \mathcal{U} of continuous and strictly monotone utility functions is:

$$\text{CCEI}_{\mathcal{U}}(\mathcal{D}) = \sup_{u \in \mathcal{U}} \min_{(p, x) \in \mathcal{D}} r(p, x; u).$$

Clearly, this coincides with Afriat's original definition when \mathcal{U} contains all continuous and monotone utility functions, but is defined given any smaller class as well. We now show a perhaps surprising, systematic relation between the CCEI and FDI (or PMI).

Proposition 7 *For any demand data \mathcal{D} and any class \mathcal{U} of continuous and strictly monotone utility function, we have: $1 - \text{CCEI}_{\mathcal{U}}(\mathcal{D}) \leq \text{PMI}_{\mathcal{U}}(\mathcal{D})$. Similarly, we have $1 - \text{CCEI}_{\mathcal{U}}(\mathcal{D}) \leq \text{FDI}_{\mathcal{U}}(\mathcal{D})$ for any class \mathcal{U} of regular utility functions.*

After reading the proof in the Appendix, it is also easy to construct examples where the inequality holds strictly, and others where it holds with equality. In the next section, we explain how any Afriat-inequality test designed to check for rationality given \mathcal{U} extends to compute the $FDI_{\mathcal{U}}$ (or $PMI_{\mathcal{U}}$). By contrast, there is no general method for computing $CCEI_{\mathcal{U}}$. Hence a side benefit of the above proposition is to provide an easy-to-compute lower bound on $CCEI_{\mathcal{U}}$.

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 WARP 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 choices 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 (or PMI) instead. Interestingly, while we have established that $1 - CCEI \leq FDI$, 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 compared to $1 - CCEI$. In the Appendix (see Example 1), we construct demand data that would appear more rational when applying Bronar's criterion to the CCEI instead of the FDI, as well as demand data that would appear more rational when applying Bronars' criterion to the FDI instead of the CCEI. Thus, when judging the quality of consumer choices, it is more prudent to assess it according to *both* measures. This is precisely what we do in Section 8, when analyzing Choi et al.'s (2014) portfolio demand data.

6 Computing the PMI and the FDI

Tests of rationalizability have been proposed over the years for many standard classes \mathcal{U} of preferences. A typical test amounts to checking whether a certain set of (often-times linear or bilinear) Afriat-like inequalities admit a solution.¹⁸ In this section, we show how all inequality-based rationalizability tests can be repurposed for computing the corresponding $\text{PMI}_{\mathcal{U}}$ and $\text{FDI}_{\mathcal{U}}$. By contrast, this is not true for the CCEI. The difference is that demand data is fully rational for modified prices under the PMI and the FDI, while consumers' choices cease to be optimal in Afriat's shrunk budget sets (indeed they are even infeasible).¹⁹

More precisely, an inequality-based test of rationalizability can be described as follows. There is a set $S_{\mathcal{U}}$ of inequalities whose coefficients are determined by the demand data, with the property that \mathcal{D} is rationalizable given \mathcal{U} if, and only if, the set of inequalities $S_{\mathcal{U}}(\mathcal{D})$ given \mathcal{D} has a solution. The next result then follows at once from the very definition of the PMI, and from the FDI-PMI coincidence identified in Proposition 2. To state it, we need some additional notation. For any demand data \mathcal{D} and consumer perceived price vectors $p^c(p, x)$ for each $(p, x) \in \mathcal{D}$, let $\mathcal{D}^c = \{(p^c(p, x), x) | (p, x) \in \mathcal{D}\}$ be the modified demand data derived from \mathcal{D} by using perceived prices instead of actual prices.

Proposition 8 *Let $S_{\mathcal{U}}$ be a test of rationalizability given \mathcal{U} . Then $\text{PMI}_{\mathcal{U}}$ is the infimum over all $\varepsilon \in [0, 1]$ such that there exists a solution to the system of inequalities listed in (3b) and in $S_{\mathcal{U}}(\mathcal{D}^c)$, with an extra variable $p^c(p, x) \in \mathbb{R}_{++}^L$ for each $(p, x) \in \mathcal{D}$. The same result applies to the FDI when all utility functions in \mathcal{U} are quasi-concave.*

For each ε , one simply checks a more permissive system of inequalities than required for the rational benchmark: the true price vector in each observation (p, x) is replaced with a perceived price vector $p^c(p, x)$, which must be close to p . Conveniently, the added inequalities stemming from (3b) are linear in the relevant exogenous

¹⁸See Diewert (2012) for an overview of Afriat-inequality tests of rationalizability applied to quasi-linear, homothetic or additively separable utility functions, and expected or non-expected utility. Polisson, Quah and Renou (2020) extend this approach to important classes of non-convex preferences, including expected utility without risk aversion.

¹⁹Alternatively, one could view a consumer's choice as optimal in the nonconvex budget set obtained by first shrinking the original budget set, and then adding back the chosen bundle. Grounded in first-order conditions, typical rationalizability tests use prices as proxies for marginal utilities (see e.g. all tests in Diewert (2012)'s survey), and cannot be adapted to such general nonconvex budget sets (in the hope of computing the CCEI). Polisson et al. (2020)'s approach can accommodate general budget sets for certain classes of preferences, and thus compute the CCEI in those cases.




Figure 2: The construction of $o(x, x')$ for proving the two-good formula.

variables (prices).²⁰ The PMI and FDI can then be identified to any desired degree of precision by following the usual dictionary-search method: first test whether the system of inequalities has a solution for $\varepsilon = 1/2$; then do the same for $\varepsilon = 1/4$ if a solution was found, or for $\varepsilon = 3/4$ otherwise; and iterate the procedure (identifying the index with $\pm \frac{1}{2^n}$ precision in n steps).

We also found a closed-form formula to compute the FDI (PMI) given the reference class \mathcal{U}^* of all regular utility functions when there are only two goods ($L = 2$). Though considering only two goods is restrictive, notice that a majority of experiments in consumer theory share that feature, and that we are not aware of a closed-form formula for the CCEI even in the two-good case. As a start, consider demand data $\mathcal{D} = \{(p, x), (p', x')\}$ comprising only two observations. There is a WARP violation if x and x' are distinct and each of these bundles is affordable when the other is chosen (i.e. $p \cdot x' \leq p \cdot x$ and $p' \cdot x \leq p' \cdot x'$). Assume $p'_1/p'_2 \geq p_1/p_2$ without loss of generality. Then,

$$\text{FDI}_{\mathcal{U}^*}(\mathcal{D}) = \begin{cases} \min\left\{\frac{p \cdot (x - x')}{p_2(x_2 - x'_2)}, \frac{p' \cdot (x' - x)}{p'_1(x'_1 - x_1)}\right\} & \text{if } \mathcal{D} \text{ violates WARP;} \\ 0 & \text{otherwise.} \end{cases}$$

²⁰All rationalizability tests we are aware of are linear in prices, and hence the tweaked system in Proposition 8 is bilinear. In many cases, including expected utility or quasi-linearity, a simple change of variable makes the system linear. Even when not reducible this way, bilinearity means that tractable or efficient computation methods are oftentimes available.

Graphically, $\text{FDI}_{\mathcal{U}^*}(\mathcal{D}) = \min\{\epsilon(p, o(x, x')), \epsilon(p', o(x, x'))\}$ where $o(x, x')$ that is orthogonal to the line passing through x and x' (see Figure 2), from which the above formula follows. To extend this result to any demand data \mathcal{D}' , we showed that, with two goods, $\text{FDI}_{\mathcal{U}^*}(\mathcal{D}')$ is the maximum of the FDI's associated to each pair of observations in \mathcal{D}' . Details to prove this appear in the earlier working paper version (de Clippel and Rozen, 2019).

7 Illustration Using Portfolio Demand Data

We now illustrate our methodology in the 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 will see how our approach is easily applied to assess decision-making quality and preference misspecification, covering classic subclasses of state-independent preferences, plain expected utility and risk-averse expected utility. Differences between the CCEI and the FDI/PMI will become apparent, but a Bronars' analysis using the FDI reinforces Choi et al. (2014)'s conclusion regarding the overall decision-making quality of their data. This section also offers further insights regarding some results appearing in an independent and contemporaneous paper on risk-averse expected utility (see Echenique, Imai and Saito (2019)).

7.1 Decision-Making Quality

We begin our analysis by focusing on the class \mathcal{U}^* of all regular utility function in this subsection. For notational simplicity, we will drop the index \mathcal{U}^* . The FDI is computed using the simple closed-form formula we provide towards the end of Section 6 for the special case of two goods. Proposition 7 tells us that $1 - \text{CCEI} \leq \text{FDI}$, and the histogram of $\text{FDI} - (1 - \text{CCEI})$ in Figure 3(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

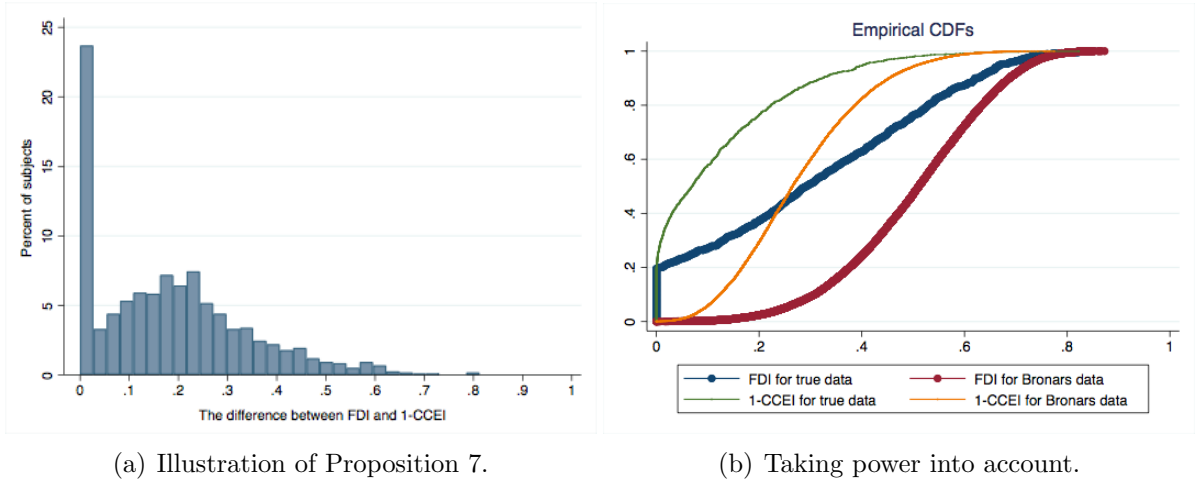


Figure 3: Illustration using data from Choi et al. (2014).

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 3(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.

7.2 FOSD and Risk-Averse Expected Utility

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 $\varepsilon/(1-\varepsilon)$). 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 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 $\text{FDI}_{\mathcal{EU}^r} \geq 1 - \text{CCEI}$.²¹

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, as seen in Figure 3(a)) has the advantage of comparing apples to apples and placing the spotlight on the dimension of interest: assessing how much more stringent expected utility is from rationality.

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.²²

²¹Translating this inequality for their parametrization gives $\text{CCEI} \geq e/(1+e)$.

²²With 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

To improve our understanding of $\text{FDI}_{\mathcal{E}U^r}$, we can decompose it as follows:²³

$$(5) \quad \text{FDI}_{\mathcal{E}U^r} = \text{FDI} + [\text{FDI}_{\mathcal{S}\mathcal{I}} - \text{FDI}] + [\text{FDI}_{\mathcal{E}U^r} - \text{FDI}_{\mathcal{S}\mathcal{I}}].$$

Thus the degree of misspecification $\text{FDI}_{\mathcal{E}U^r} - \text{FDI}$ 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.

One can show 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

$$(6) \quad \overline{\mathcal{D}} = \mathcal{D} \cup \{((p_2, p_1), (x_2, x_1)) \mid (p, x) \in \mathcal{D}\}$$

is ε -rationalizable by a regular preference.²⁴ Thanks to this observation, we can easily compute $\text{FDI}_{\mathcal{S}\mathcal{I}}$ by applying once again the simple closed-form formula provided towards the end of Section 6.

Figure 4(a) depicts the empirical CDF's of the FDI, $\text{FDI}_{\mathcal{S}\mathcal{I}}$, and $\text{FDI}_{\mathcal{E}U^r}$ for the data,²⁵ where $\text{FDI}_{\mathcal{E}U^r}$ is computed using Proposition 8, 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 $\text{FDI}_{\mathcal{E}U^r} \approx \text{FDI}_{\mathcal{S}\mathcal{I}} > \text{FDI}$. For each subject, consider

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 $\text{FDI}_{\mathcal{E}U^r}$.

²³See Halevy et al. (2018) for similar decompositions with the CCEI for parametric classes.

²⁴Necessity 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 key step of the proof is then to show the existence of such a utility function that ε -rationalizes $\overline{\mathcal{D}}$, see earlier working paper version (de Clippel and Rozen (2019)) for details.

²⁵Polisson et al. (2020, Figure 7) use their GRID method to construct an analogue of Figure 4(a) for the CCEI.

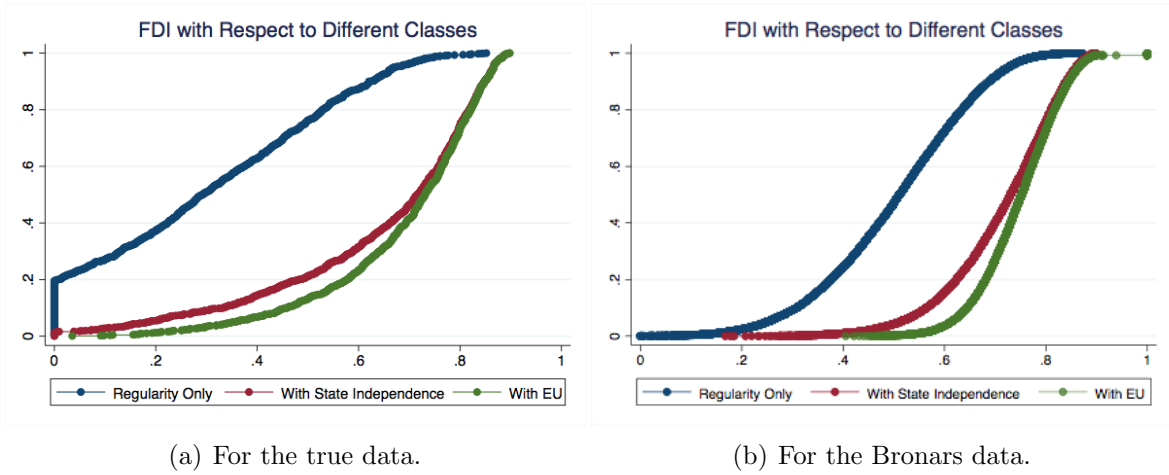


Figure 4: Empirical CDF's of the FDI with respect to the class of regular utility functions and the subclasses of state-independent utility \mathcal{SI} and risk-averse expected utility \mathcal{EU}^r .

the EU-misspecification ratio:

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

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 5 suggests the empirical CDF of the misspecification ratio for the true data is not much better than that for the randomly-generated, Bronars data.²⁶

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 budget set $p_1x_1 + p_2x_2 = 1$, with $p_1 > p_2$. The probability that she appears rational (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.

²⁶For comparability, we use the same, realized Bronars dataset throughout.




Figure 5: Empirical CDFs of the EU-misspecification ratio.

Of course, the experiment tests subjects in many more budget sets, which makes the comparison more complex. Figure 4(b) depicts the same objects as in Figure 4(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.

7.3 Plain Expected Utility

We just found out that the empirical departure from risk-averse expected utility maximization is mostly attributable to departures from plain preference maximization with regular utility functions that are state independent, and that only richer data could disentangle risk-averse expected utility on top of state independence. Did quasi-concavity play an important role in reaching this conclusion? While there is not evidence of risk-loving behavior in the data,²⁷ accommodating non-convex preferences

²⁷A 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

remains important. For instance, considering a parametric class of Gul (1991) preferences capturing disappointment aversion, Halevy et al. (2018) finds in a similar dataset that some subjects are best described via a negative parameter of disappointment aversion (in which case they are elation loving). Such preferences are state independent, but neither convex nor concave.

We can assess the impact of quasi-concavity in our analysis of the Choi et al. data by investigating how the curves in the two panels of Figure 4 move when applying the PMI (instead of the FDI), after dropping quasi-concavity in each reference class of utility functions. First, it follows from Proposition 3 that the top curves (blue) remain unchanged. Second, we show in the Appendix (Proposition 9) that dropping quasi-concavity makes no difference for state-independent utility functions either: the middle curve (red) remains unchanged as well. Third, dropping quasi-concavity can make a difference when it comes to expected utility.²⁸ But of course, the bottom curve (green) cannot move above the middle one (red) since expected utility is more demanding than state independence. At the same time, the green curve can only move up: adding risk aversion on top of expected utility increase the PMI, and the PMI coincides with the FDI on any class of regular utility functions (including the class of risk-averse expected utility). Thus our view of the data from Section 7.2 remains valid in the absence of quasi-concavity.

8 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 theo-

in more than 90% of the budget sets they faced.

²⁸See Section A4.1 in the online appendix of Polisson et al. (2020). As discussed earlier, their test can be repurposed along the lines of Proposition 8 to compute the PMI given risk-averse expected utility (both for the true data and the randomly-generated Bronars choices). We did not perform these computations, as we are proving in this paragraph that the PMI’s CDF must be squeezed between the green and red curves in Figure 4, which happen to be very close to each other.




Figure 6: Illustration of the taxation example.

ries. These studies illustrate the importance of relaxing the assumption of perfect collinearity between marginal prices and the utility gradient.

Consider for instance labor supply decision with progressive taxation. There are two goods: good 1 corresponds to leisure (as opposed to work); good 2 represents consumption. Normalizing the price of consumption to 1, the price of leisure corresponds to the real hourly wage. The consumer is endowed with H waking hours (for instance, over the course of a year), and decides how many hours to work. Progressive taxation means that there is a succession of income thresholds $0 < w_1 < w_2 < \dots$ with an increasing marginal tax rate $t_1 < t_2 < \dots$ applying to successive income brackets (t_i applying to the bracket $[w_{i-1}, w_i]$). Thus the tax on a yearly wage $W \in [w_{k-1}, w_k]$ is $t_k(W - w_{k-1}) + \sum_{j=1}^{k-1} t_j(w_{j-1} - w_j)$. Clearly, this leads to convex budget sets with piecewise-linear frontiers, which will change as the tax code and hourly wages vary.

Figure 6 provides a simple illustration. With $H = 4,000$ for a year, the consumer 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 $66\frac{2}{3}\%$. Selecting $x = (3K, 20K)$ in year 1 and $x' = (1.5K, 37.5K)$ in year 2 is inconsistent with rationality (a WARP violation).

How does one measure the extent of departure from rationality? It is unclear what the CCEI might be in this case, as there are multiple conflicting ways to ‘shrink’ 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. For instance, 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 endowment 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 with non-linear prices. By contrast, the notion of the FDI straightforwardly applies to all problems where choices are characterized by a first-order condition. With progressive taxation as above, one simply measures the discrepancy between the utility gradient and the *marginal* price vector at chosen bundles. It is not difficult to check that the FDI for the above demand data is then roughly 28.5% (details available in our earlier working paper version, see de Clippel and Rozen (2019)).

Given the prevalence of first-order conditions, the FDI can shed light on decision-making quality and misspecification in numerous other situations in economics and game theory. Here are just a few examples. One may be interested, for instance, in better understanding people’s rationality and preferences (regarding say altruism, risk or time) by collecting data in experiments with general convex budget sets instead of restricting attention to linear ones, which have been the norm so far.

One may also be interested in producers instead of consumers. Consider a firm selling its output in a competitive market, and deciding its employees’ work hours. The profit maximization benchmark corresponds to the following maximization problem:

$$\max_{L \geq 0} [pf(L) - W(L)],$$

Demand data $\{(p_k, W_k, L_k) \mid k = 1, \dots, K\}$ comprises combinations of exogenous output price (p_k), exogenous wage function (W_k), and endogenous labor demand (L_k). The wage could be linear: $W_k(L) = w_k L$, for some observable hourly wage w_k . But more realistically, the firm may have to factor in overtime wages, leading to a piecewise-linear convex function W_k . The production function f is unknown, and simply assumed to be strictly increasing, concave and differentiable (non-parametric

approach). ε -Rationalizability given this benchmark model is simply characterized by the relaxed first-order conditions:

$$1 - \varepsilon \leq \frac{p_k f'(L_k)}{W'_k(L_k)} \leq \frac{1}{1 - \varepsilon}.$$

The FDI is the infimum over all ε for which one can *find* a production function f satisfying the above inequalities for all k . As in the rest of the paper, solving this non-parametric problem may seem daunting at first since the infimum is computed over a large class of functions. But, as often happens (see Section 6), the problem reduces to a linear-programming problem. Indeed, we just need to find marginal product parameters $\beta_k = f'(L_k)$, with the key property that they are decreasing in the labor input. It is straightforward to check²⁹ that the FDI is the infimum of $1 - \min_{k=1}^K \left\{ \frac{p_k \beta_k}{W'_k(L_k)}, \frac{W'_k(L_k)}{p_k \beta_k} \right\}$ over all vectors $\beta \in \mathbb{R}_{++}^K$ such that (a) $\beta_i > \beta_j$ if $L_i < L_j$, and (b) $\beta_i = \beta_j$ if $L_i = L_j$ (for all $1 \leq i, j \leq K$). For instance, the data is consistent with profit maximization ($FDI = 0$) if, and only if, (a') $W'_i(L_i)/p_i > W'_j(L_j)/p_j$ if $L_i < L_j$, and (b') $W'_i(L_i)/p_i = W'_j(L_j)/p_j$ if $L_i = L_j$. Indeed, β_k must equal $W'_k(L_k)/p_k$ in that case. Violations of this co-monotonicity property leads to larger FDIs, letting β take a wider range of values in order to satisfy (a) and (b) above.

In some cases, the modeler observes output, but not labor. Even when labor is also observed, the modeler might not observe the wage function. Our methodology applies here too. Demand data $\{(p_k, y_k) \mid k = 1, \dots, K\}$ comprises combinations of exogenous output price (p_k), and endogenous production level (y_k). As before, the production function f is unknown, and simply assumed to be strictly increasing, concave and differentiable. But now the wage function W is also unknown, and assumed to be strictly increasing, convex and differentiable. ε -Rationalizability given this benchmark model is simply characterized by the relaxed first-order conditions:

$$1 - \varepsilon \leq \frac{p_k f'(f^{-1}(y_k))}{W'(f^{-1}(y_k))} \leq \frac{1}{1 - \varepsilon}.$$

The FDI is the infimum over all ε for which one can *find* a production function f and a wage function W satisfying the above inequalities for all k . We can simplify this once again, to finding parameters $\gamma_k = \frac{f'(f^{-1}(y_k))}{W'(f^{-1}(y_k))}$ satisfying the inequalities and

²⁹The reasoning in the text establishes necessity. For sufficiency, one simply constructs a valid function f whose derivative at L_k is β_k , which is always possible if the vector β satisfies (a) and (b).

decreasing in y_k : as output and thus labor increases, the marginal product of labor decreases while the marginal wage increases. It is then easy to check that the FDI is the infimum of $1 - \min_{k=1}^K \{p_k \gamma_k, \frac{1}{p_k \gamma_k}\}$ over all vectors $\gamma \in \mathbb{R}_{++}^K$ such that (a) $\gamma_i > \gamma_j$ if $y_i < y_j$, and (b) $\gamma_i = \gamma_j$ if $y_i = y_j$ (for all $1 \leq i, j \leq K$). For instance, the data is consistent with profit maximization ($FDI = 0$) if, and only if, (a') $p_j > p_i$ if $y_j > y_i$, and (b') $p_i = p_j$ if $y_i = y_j$. Violations of this co-monotonicity property leads to larger FDIs, letting γ take a wider range of values in order to satisfy (a) and (b) above.

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.
- BALAKRISHNAN, U. AND J. HAUSHOFER AND P. JAKIELA (2019), Evidence of Present Bias from Convex Time Budget Experiments, *Experimental Economics*, forthcoming.
- 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, J., T. HOSSAIN AND J. MORGAN (2010), Shrouded Attributes and Information Suppression: Evidence from the Field, *Quarterly Journal of Economics*, 125(2), 859–876.
- CHETTY, R., A. LOONEY AND K. KROFT (2009), Salience and Taxation: Theory and Evidence, *American Economic Review*, 99, 1145–1177.
- 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 (2019), Relaxed Optimization: ε -Rationalizability and the FOC-Departure Index in Consumer Theory, *SSRN working paper*.

- 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.
- FISMAN, R. AND S. KARIV AND D. MARKOVITS (2007), Individual Preferences for Giving, *American Economic Review*, 97(5), 1858–1876.
- FRISBY, J. AND J. STONE (2010), *Seeing*, Massachusetts Institute of Technology.
- 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.
- HOUTMAN, M. AND J.A.H. MAKS (1985), Determining all Maximal Data Subsets Consistent with Revealed Preference, *Kwantitatieve Methoden*, 19, 89–104.
- ITO, K. (2014), Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing, *American Economic Review*, 104(2), 537–563.
- KAHNEMAN, D., AND A. TVERSKY (1979), An Analysis of Decision under Risk, *Econometrica*, 47(2), 263–291.
- KÖSZEGI, B., M. RABIN (2006), A Model of Reference-Dependent Preferences, *The Quarterly Journal of Economics*, 121(4), 1133–1165.
- POLISSON, M., J. QUAH AND L. RENO (2020), Revealed Preferences over Risk and Uncertainty, *American Economic Review*, 110, 1782–1820.

- 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.
- THOMAS, M. AND V. MORWITZ (2005), Penny Wise and Pound Foolish: The Left-Digit Effect in Price Cognition, *Journal of Consumer Research*, 32(1), 54–64.
- VARIAN, H.R. (1982), The Nonparametric Approach to Demand Analysis, *Econometrica*, 50(4), 945–973.
- TVERSKY, A. AND D. KAHNEMAN (1991), Loss Aversion in Riskless Choice: A Reference-Dependent Model, *The Quarterly Journal of Economics*, 106(4), 1039–1061.
- VARIAN, H.R. (1990), Goodness-of-Fit in Optimizing Models, *Journal of Econometrics*, 46, 125–140.

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

$$(x, y) \succeq (x', y') \Leftrightarrow \left(\left(\frac{x_1/x_2}{y_1/y_2}, 1 \right), (1, 1) \right) \succeq \left(\left(\frac{x'_1/x'_2}{y'_1/y'_2}, 1 \right), (1, 1) \right).$$

Representability implies the existence of a function $f : \mathbb{R}_{++}^2 \times \mathbb{R}_{++}^2 \rightarrow \mathbb{R}$ such that $(x, y) \succeq (x', y')$ if and only if $f(x, y) \geq f(x', y')$. Defining $g(\gamma) := f((\gamma, 1), (1, 1))$, 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 $\gamma \geq 1$. By Numbering and Unit Invariance, $g(\gamma) = g(1/\gamma)$ and thus $g(\gamma) = g(\max\{\gamma, 1/\gamma\})$. As $\max\{\gamma, 1/\gamma\} \geq 1$, inverting g gives the result. *Q.E.D.*

Proof of Proposition 5 (ε -Rationalizability and additively separable utility)

A preliminary observation will prove useful. Fix $\beta \in \mathbb{R}_{++}^L$. As easily checked, $x \in \mathbb{R}_+^L$ maximizes $\sum_{\ell} \beta_{\ell} u_{\ell}(y_{\ell})$ over the set of bundles $y \in \mathbb{R}_+^L$ such that $p \cdot y \leq p \cdot x$ if, and only if, it maximizes $\sum_{\ell} u_{\ell}(y_{\ell})$ over the set of bundles $y \in \mathbb{R}_+^L$ such that $\frac{p}{\beta} \cdot y \leq \frac{p}{\beta} \cdot x$, where $\frac{p}{\beta}$ is the vector defined by $(\frac{p}{\beta})_{\ell} = p_{\ell}/\beta_{\ell}$ for each ℓ .

Fix $(p, x) \in \mathcal{D}$. Remember from the proof of Proposition 2 that u ε -rationalizes (p, x) if, and only if, there exists $p^c \in \mathbb{R}_{++}^L$ satisfying (3b) such that x maximizes $\sum_{\ell} u_{\ell}(y_{\ell})$ over the set of bundles $y \in \mathbb{R}_+^L$ such that $p^c \cdot y \leq p^c \cdot x$. The result then follows from the above observation by defining β by $\beta_{\ell} = \frac{p_{\ell}}{p_{\ell}^c}$ (so that $\frac{p}{\beta} = p^c$). *Q.E.D.*

Proof of Proposition 6 (Reference Dependence)

Notice that x maximizes $v(y|r(p, x))$ over the set of bundles $y \in \mathbb{R}_+^L$ such that $p \cdot y \leq p \cdot x$ if, and only if, there exists $\beta(p, x) \in \mathbb{R}_{++}^L$ such that x maximizes $\sum_{\ell=1}^L \beta_{\ell}(p, x) u_{\ell}(y)$ over the same set of bundles, $\beta_{\ell}(p, x) = 1 + \gamma$ if $x_{\ell} > r_{\ell}(p, x)$, $\beta_{\ell}(p, x) = 1 + \gamma\lambda$ if $x_{\ell} < r_{\ell}(p, x)$, and $\beta_{\ell}(p, x) \in [1 + \gamma, 1 + \gamma\lambda]$ if $x_{\ell} = r_{\ell}(p, x)$.

For necessity, suppose that \mathcal{D} is consistent with reference-dependent choices for gain-loss parameters $\gamma \geq 0$ and $\lambda \geq 1$. By Proposition 5, $\sum_{\ell} u_{\ell}$ ε -rationalizes \mathcal{D} for ε equal to 1 minus the smallest ratio of $\beta_{\ell}(p, x)/\beta_{\ell'}(p, x)$ over all ℓ, ℓ' , and $(p, x) \in \mathcal{D}$. This minimum is at least $\frac{1+\gamma}{1+\gamma\lambda}$, and hence ε is lower or equal to $\frac{\gamma(\lambda-1)}{1+\gamma\lambda}$, as desired.

For sufficiency, suppose that \mathcal{D} is ε -rationalizable given $\sum_{\ell} u_{\ell}$. By Proposition 5, for all $(p, x) \in \mathcal{D}$, there is $\beta(p, x) \in \mathbb{R}_{++}^L$ such that (i) x maximizes $\sum_{\ell=1}^L \beta_{\ell}(p, x) u_{\ell}(y)$ over the set of bundles $y \in \mathbb{R}_+^L$ such that $p \cdot y \leq p \cdot x$, and (ii) $\frac{1+\gamma}{1+\gamma\lambda} \leq \frac{\beta_{\ell}(p, x)}{\beta_{\ell'}(p, x)} \leq \frac{1+\gamma\lambda}{1+\gamma}$, for all $\ell, \ell' \in \{1, \dots, L\}$. Renormalizing $\beta(p, x)$, if needed,³⁰ we have that $\beta(p, x)$'s components are all between $1 + \gamma$ and $1 + \gamma\lambda$. By the first paragraph of this proof, \mathcal{D} is consistent with reference-dependent choices for gain-loss parameters $\gamma \geq 0$ and $\lambda \geq 1$, simply by taking $r(p, x) = x$ for all $(p, x) \in \mathcal{D}$. *Q.E.D.*

Proof of Proposition 7 (Relation to CCEI)

Take any $\varepsilon > \text{PMI}_{\mathcal{U}}(\mathcal{D})$. One can find $u \in \mathcal{U}$ and associate to all $(p, x) \in \mathcal{D}$ a price vector $p^c(p, x) \in \mathbb{R}_{++}^L$ such that conditions (3a) and (3b) hold. Take any $(p, x) \in \mathcal{D}$, and consider the indifference curve of u passing through x . As illustrated in Figure 7

³⁰Multiplying $\beta(p, x)$ by $(1 + \gamma)$ over $\beta(p, x)$'s smallest component, which does not change (i).




Figure 7: Construction for the proof of Proposition 7.

(see the dashed line), (3a) means the perceived budget line associated to $p^c(p, x)$ (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 this hyperplane, i.e., any y such that $p^c(p, x) \cdot y \geq p^c(p, x) \cdot x$. This holds trivially if $p^c(p, x) = p$, so suppose they are different and consider the optimization problem

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

The constraint must bind at the optimum, else the objective could be further reduced. As seen in Figure 7, linearity of the objective and constraint imply the objective function is minimized at a bundle y with only one positive component: that is, there is ℓ such that $y_\ell = \frac{p^c(p, x) \cdot x}{p_\ell^c(p, x)}$ and $y_i = 0$ for all $i \neq \ell$. Using (3b), the minimal expenditure satisfies:

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

By quasi-concavity of u , any bundle z with $u(z) \geq u(x)$ must satisfy $p^c(p, x) \cdot z \geq p^c(p, x) \cdot x$. 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. Hence $CCEI_{\mathcal{U}}(\mathcal{D}) \geq 1 - \varepsilon$. This proves the result since the inequality holds for all $\varepsilon > \text{PMI}_{\mathcal{U}}(\mathcal{D})$. *Q.E.D.*

Example 1 (Data Not Systematically More Rational under CCEI or FDI) *Consider demand data $\mathcal{D} = \{(p, x), (p', x')\}$ with two observations and a WARP 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 (in this example we focus on the reference class of all regular utility functions), 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 WARP violation: $p \cdot y' \leq p \cdot y = p \cdot x$ and $p' \cdot y \leq p' \cdot y' = p' \cdot x'$. By the simple closed-form formula given at the end of Section 6, the FDI of such choices is

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

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}' .³¹ 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}' .

Proposition 9 Consider any portfolio demand data \mathcal{D} as in Section 7. Let \mathcal{SI}^- be the class of utility functions derived from \mathcal{SI} by dropping quasi-concavity. Then $\text{PMI}_{\mathcal{SI}^-}(\mathcal{D}) = \text{FDI}_{\mathcal{SI}}(\mathcal{D})$.

Proof. Note that: $\text{FDI}_{\mathcal{SI}}(\mathcal{D}) = \text{FDI}_{\mathcal{U}}(\bar{\mathcal{D}}) = \text{PMI}_{\mathcal{U}^-}(\bar{\mathcal{D}}) = \text{PMI}_{\mathcal{SI}^-}(\mathcal{D})$, where $\bar{\mathcal{D}}$ is the mirror-extended dataset from (6), \mathcal{U} is the class of regular utility functions, and \mathcal{U}^- is

³¹For 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 $\text{FDI}(\{(p, \hat{x}), (p', \hat{z})\}) > \text{FDI}(\mathcal{D}')$.

derived from \mathcal{U} by dropping quasi-concavity. The first equality was already mentioned earlier (see footnote 24), and the second was established in Proposition 3. 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 $\text{PMI}(\bar{\mathcal{D}}) \leq \text{PMI}_{\mathcal{SI}^-}(\mathcal{D})$. 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 $x_2 \geq x_1$. *Q.E.D.*