ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

De Clippel, Geoffroy; Rozen, Kareen

Working Paper

Consumer theory with misperceived tastes

Working Paper, No. 2018-10

Provided in Cooperation with:
Department of Economics, Brown University

Suggested Citation: De Clippel, Geoffroy; Rozen, Kareen (2018) : Consumer theory with misperceived tastes, Working Paper, No. 2018-10, Brown University, Department of Economics, Providence, RI

This Version is available at:
https://hdl.handle.net/10419/202602

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

CONSUMER THEORY WITH MISPERCEIVED TASTES

Geoffroy de Clippel* and Kareen Rozen†

July 2018

Abstract

Incorporating bounded rationality into the classic consumer theory setting, we study the testable implications of a consumer who may have trouble consistently assessing her subjective tastes. Our model of ε-Rationalizability, which bounds the consumer's misperception of her marginal rates of substitution, may arise from various choice heuristics. It also offers a natural, preference-based measure of departure from rationality that is more demanding than Afriat's measure.

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

1. Introduction

Recent years have seen a surge of interest in understanding bounded rationality, with much of this work concentrating on discrete choice domains. In this paper, we extend the classic consumer theory setting (Afriat, 1967; Varian, 1982; Chiappori and Rochet, 1987; Matzkin and Richter, 1991) to accommodate bounded rationality. Specifically, we study a consumer who may misperceive her utility tradeoffs.

A building block of our approach is that the consumer equates her *perceived* marginal rates of substitution with opportunity costs. Otherwise, as classic consumer theory tells us for the case of accurate perceptions, the consumer would wish to adjust her bundle in some direction. Given a utility function u, a consumer's marginal rate of substitution, $MRS^u_{\ell\ell'}(x)$, captures the local utility tradeoffs she faces between two goods ℓ and ℓ' as she shifts away from a bundle x. Our consumer may misperceive her true marginal rates of substitution to a limited extent, and thus deviate from rationality. As formalized in Section 2, we say that demand data is ε-rationalizable if there exists a utility function u, satisfying some regularity properties,[1] such that whenever we observe the bundle x being chosen at a price vector p, we have

$$1-\varepsilon \leq \frac{MRS^u_{\ell\ell'}(x)}{p_\ell/p_{\ell'}} \leq \frac{1}{1-\varepsilon}, \text{ for all } \ell \neq \ell'.$$

Marginal rates of substitution (MRS) and opportunity costs must match when $\varepsilon = 0$, as must occur under rational choice. The larger ε is, the more permissive ε-Rationality becomes. Importantly, the theory is invariant to the units in which goods are measured and priced, as it quantifies the departure between MRS and opportunity costs in percentage terms. Moreover, it treats overestimation and underestimation of MRS in a symmetric fashion.

One can think of various heuristics leading to ε-rationalizable demand data. Instead of testing a tentative choice against all affordable alternatives, a consumer may perform only local comparisons. Such limited attention[2] is rational when preferences are convex, and even more sensible than considering distant alternatives when con-

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

[2] Models of limited attention have been studied recently; see, among other works, attention grabbing in Eliaz and Spiegler (2011); shortlisting in Manzini and Mariotti (2007); categorization in Manzini and Mariotti (2012); forms of unawareness in Masatlioglu, Nakajima and Ozbay (2012); rationalization in Cherepanov, Feddersen and Sandroni (2013); optimally sparse attention in Gabaix (2014); and choice overload in Lleras, Masatlioglu and Nakajima (2017).

templation is costly. At the same time, it is plausible that the consumer's perception[3] of tradeoffs at the margin is imperfect, resulting in ε-rationalizable choices. Many factors, some unobservable to the data analyst, could potentially impact the consumer's perception of her tradeoffs. The consumer may observe another customers' choices; notice unsatisfactory expiration dates; interact with a persuasive salesperson; or be swayed by visuals, such as store displays and packaging. Furthermore, the initial bundle she contemplates may affect the resting point of dynamics in which she compares prices with perceived utility tradeoffs. As a result, the consumer might be observed demanding different bundles on different occasions of facing the same price vector. Our framework accommodates the possibility of such data.

The notion of ε-Rationalizability also relates to another strand of the bounded rationality literature. Kalai, Rubinstein and Spiegler (2002) suggest that the decision maker does maximize a preference, but that it may vary with the choice problem she faces. One can also think of more general forms of context-dependence, where a decision maker maximizes different preferences when facing the same problem multiple times. While such theories lack empirical content in general, nontrivial testable implications may arise if one restricts the set of acceptable preferences to 'distortions' of an underlying welfare preference, as suggested by Rubinstein and Salant (2012). As we later formalize, ε-Rationalizability is behaviorally equivalent to a theory in which the DM uses only preferences sufficiently close to her true utility function u, where distance is measured using percentage change in MRS. This has an especially intuitive interpretation in the additively-separable case (studied in Section 4.3), where we show that the nearby utility functions differ only in the weights of the component utilities being added. For instance, this means that when choosing a consumption stream, any misperception in MRS can be attributed to misperception of discounting functions by a scaling factor, in the spirit of quasi-hyperbolic discounting.

We analyze the testable implications of ε-Rationalizability in Section 3, studying under what conditions there exists a regular utility function such that each datapoint corresponds to at most an ε-misperception of marginal rates of substitution. If the data is ε-rationalizable, then each demanded bundle x can be associated to some vector v in a cone $C_\varepsilon(x)$ representing possible misperceived rates of substitution around

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

x. While the modeler cannot conclude that a chosen bundle x is revealed preferred to all other feasible ones, she can use quasi-concavity to conclude that x is revealed preferred to all the bundles from which it is separated by some vector $v \in C_\varepsilon(x)$. The difficulty is that the modeler does not know which $v \in C_\varepsilon(x)$ is the true gradient of u at x, and thus which v restricts the lower-contour set of x under the consumer's true preference. This can be addressed by the techniques of de Clippel and Rozen (2018), who generalize the approach for testing SARP (the Strong Axiom of Revealed Preference) to bounded rationality models. While the consumer's utility function is defined over the entire domain of bundles $\mathbb{R}^L_+$, we show that testing ε-Rationalizability is equivalent to checking whether there exists an acyclic preference relation satisfying these restrictions over the *finite* subset of demanded bundles; and that testing is tractable using de Clippel and Rozen (2018)'s enumeration procedure. Indeed, when the test succeeds, then we can use the misperceived gradients it constructs to make an auxiliary rational dataset. Classic results (Chiappori and Rochet, 1987; Matzkin and Richter, 1991) then apply to generate a utility function for the consumer.

We further develop the testable implications in Section 4, where we study perhaps the three most classic restrictions on consumer utility functions: quasi-linearity, homotheticity, and additively separability. Classic testing results for rationalizability under each of these assumptions look for the existence of a solution to a system of inequalities. This section shows how these results generalize to ε-Rationalizability, providing graphical intuition in the quasi-linear and homothetic cases, and shedding light on sources of misperception of MRS in the additively-separable case.

Our model naturally lends itself as a measure of choice inconsistency. In Section 5, we define the Tradeoff-Misperception Index of demand data as the minimal ε that makes it ε-rationalizable. Thus, the Tradeoff-Misperception Index measures the extent of local errors in assessing marginal rates of substitution. By contrast, the classic Critical Cost Inefficiency Index of Afriat (1973) considers global monetary effects, measuring the percentage of income lost due to irrationality. Despite these differences, there is a surprising relationship between the two measures. It turns out that our measure is more demanding, as a small local error in perception implies only small budgetary adjustments are needed to eliminate revealed preference cycles. Formally, we show that whenever the Tradeoff-Misperception Index is smaller than ε, then so is Afriat's index, *but not vice-versa*. We also show that there exist datasets where the two indices coincide. Nonetheless, there are some benefits and

drawbacks to each index. Our measure requires a stronger notion of preference regularity than Afriat's; but as we discuss later, our measure is also more robust to non-linear budgets and other choice domains, as it involves no adjustments to constraint sets. Finally, following the approach of Halevy, Persitz and Zrill (2018), we observe that the Tradeoff-Misperception Index can be used to study goodness-of-fit of parametric classes of utility functions.

There are other possible interpretations and applications of our model. When equating marginal rates of substitution and opportunity costs, one might imagine that misperception applies to prices instead of tastes. We conclude the paper by considering this interpretation in Section 6, and discuss connections to Gabaix (2014).

2. Consumer Data and ε-Rationalizability

We observe a consumer selecting a consumption bundle at various price vectors. The *demand data* $\mathcal{D}$ comprises a finite collection of pairs (p, x), where $p \in \mathbb{R}^L_{++}$ is a *price vector* and $x \in \mathbb{R}^L_{++}$ is the *consumption bundle* demanded at p. Note that a price vector, or a demanded bundle, may potentially appear in $\mathcal{D}$ multiple times.[4]

The rational benchmark posits that the consumer selects bundles through utility maximization over the budget set, whereby opportunity costs and marginal rates of substitution are equalized. Doing this correctly requires the consumer to have an accurate understanding of both prices and her utility function. We primarily[5] consider settings where consumers do assess prices accurately, but may find it difficult to consistently assess their subjective tastes.

Say that a utility function is *regular* if it is differentiable, strictly monotone and strictly quasi-concave. In that case, the *marginal rates of substitution* (MRS) associated to any two goods is strictly positive and well defined by:

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

for all $\ell \neq \ell'$, at the bundle x. By contrast, the opportunity cost between goods ℓ and ℓ' is given by the price ratio $p_\ell/p_{\ell'}$.

[4] As observed in Chiappori and Rochet (1987), consistency with rationality using a differentiable utility function would rule out the same bundle to be chosen at multiple price vectors. With misperception, however, this may occur.

[5] In Section 6, we explore misperceived prices and draw connections with our main results.

Definition 1 (ε-Rationalizability) *The demand data $\mathcal{D}$ is ε-rationalizable if there exists a regular utility function u such that*

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

for each $\ell \neq \ell'$ and each $(p,x) \in \mathcal{D}$. Then u is said to ε-rationalize *the demand data.*

One way to think of an ε-rationalizable consumer is to imagine her performing only local comparisons: would small shifts away from the bundle in question improve utility? This is the calculation performed in classic consumer theory when equating opportunity costs and marginal rates of substitution, and results in a globally optimal choice when preferences are convex (as assumed here). The critical difference is that our consumer may have trouble consistently assessing her utility tradeoffs at the margin, leading to choices inconsistent with rationality. Among other causes, such misperception could be due to the mere subjectivity of utility (thus leading to some natural fluctuations of assessments), or unobservable factors affecting the appeal of products at the time of purchase.

The notion of ε-Rationalizability also has an interesting connection with Rubinstein and Salant (2012), who suggest that welfare analysis may be performed using irrational data if we can find a true preference relation such that each of the DM's choices can be explain by maximizing some preference that is not too far from the true one in some sense (e.g., flipping at most a couple of comparisons). In our setting, we can use ε-Rationalizability to define a notion of distance. Say that two regular utility functions u and u' are within distance ε if $MRS^{\hat{u}}_{\ell\ell'}/MRS^u_{\ell\ell'} \in [1-\varepsilon, \frac{1}{1-\varepsilon}]$ for all $x \in \mathbb{R}^L_+$ and all $\ell \neq \ell' \in \{1,\ldots,L\}$. Then, $\mathcal{D}$ is ε-rationalizable if and only if there is a regular utility function u such that for all $(p,x) \in \mathcal{D}$, x is $\hat{u}$-maximal among bundles costing at most $p \cdot x$, for some regular utility function $\hat{u}$ within ε distance of u.[6]

3. Testable Implications

The seminal work of Afriat (1967) shows how the generalized axiom of revealed preference (GARP) captures the empirical content of rational choice. Formally, demand data is consistent with the maximization of some strongly monotone, continuous

[6]Sufficiency is trivial. For necessity, take any $(p,x) \in \mathcal{D}$ and define $\hat{u}(x'_1,\ldots,x'_L) := u(\delta_1 x'_1,\ldots,\delta_L x'_L)$ $\forall x' \in \mathbb{R}^L_+$, where $\delta_\ell := \frac{p_\ell}{\partial u(x)/\partial x_\ell}$. Clearly, $\hat{u}$ inherits strict monotonicity and differentiability from u, and $\hat{u}$ must be within ε distance of u. It is not difficult to show that $\hat{u}$ also inherits strict quasi-concavity. Then x must also be $\hat{u}$-optimal, as it satisfies the FOC.

utility function if and only it satisfies GARP. Given the widespread use of regularity, subsequent work by Chiappori and Rochet (1987) and Matzkin and Richter (1991) extended Afriat's work in that direction. They show that only a slightly stronger requirement than GARP arises.[7]

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

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

$$\frac{p_\ell}{p_{\ell'}}(1-\epsilon) \leq MRS^u_{\ell\ell'}(x) \leq \frac{p_\ell}{p_{\ell'}}\left(\frac{1}{1-\epsilon}\right), \quad \forall\, \ell \neq \ell'. \tag{1}$$

The linear inequalities in (1) define a cone $C_\varepsilon(p, x) \subseteq \mathbb{R}^L_{++}$ to which the gradient of u at x must belong. We must keep in mind, however, that the bundle x is potentially demanded under multiple price vectors. Let $P(x) = \{p' \in \mathcal{R}^L_{++} \mid (p', x) \in \mathcal{D}\}$ be the set of price vectors at which x is demanded. To capture the bounds imposed jointly by the data on the gradient of the utility function at x, we must define the cone:

$$C_\varepsilon(x) = \bigcap_{p' \in P(x)} C_\varepsilon(p', x). \tag{2}$$

Suppose we conjecture that the gradient of u at x is given by the vector $v \in C_\varepsilon(x)$. As u is strictly quasi-concave, $u(x) > u(x')$ for any bundle $x' \neq x$ such that $v \cdot x' \leq v \cdot x$. In particular, this implies x is strictly preferred to the subset of such x' that were also observed chosen, namely, to bundles in the set:

$$\Gamma(x, v) = \{x' \in X \mid x' \neq x,\ v \cdot x' \leq v \cdot x\}. \tag{3}$$

[7]Strict quasi-concavity implies the demand function is single-valued. In particular, a chosen bundle is revealed *strictly* preferred to all other affordable alternatives, so the data satisfies SARP, not just GARP. Requiring differentiability adds an 'invertibility' requirement that the consumer cannot pick the same bundle from two different budget sets.

The difficulty, of course, is that we are unsure which $v \in C_\varepsilon(x)$ is the true gradient of u at x. We only know that there *exists* $v \in C_\varepsilon(x)$ such that x is strictly preferred to all the elements in $\Gamma(x, v)$. In the language of de Clippel and Rozen (2018), this corresponds to a *lower-contour set restriction* on the consumer's preference $\succ$ over the set X of demanded bundles, since some element of $\{\Gamma(x, v) \mid v \in C_\varepsilon(x)\}$ must be contained in the $\succ$-lower contour set of x. There is one such lower-contour set restriction for *each* $x \in X$, generating a collection of lower-contour set restrictions $\mathcal{R}_\varepsilon(\mathcal{D})$ over the consumer's possible preference over demanded bundles. Thus, we have shown that ε-Rationalizability of the demand data $\mathcal{D}$ requires the collection of restrictions $\mathcal{R}_\varepsilon(\mathcal{D})$ to be *acyclically satisfiable*: there must exist a (strict) acyclic relation on X that simultaneously satisfies them. For instance, the consumer's true utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$, restricted to X, defines such an acyclic relation.

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

Theorem 1 *Acyclic satisfiability of $\mathcal{R}_\varepsilon(\mathcal{D})$ is necessary and sufficient for the demand data $\mathcal{D}$ to be ε-rationalizable.*

Proof. It remains to prove sufficiency. By acyclic satisfiability, there is a strict ordering $\succ$ over X with the feature that for all $x \in X$, there is $v(x) \in C_\varepsilon(x)$ such that $x \succ x'$ for all $x' \in \Gamma(x, v)$. We now construct an auxiliary demand data $\mathcal{D}' = (v(x), x)_{x \in X}$ where the vector $v(x)$, which is strictly positive, is taken to be the price vector when x is chosen. This data satisfies SARP, since any cycle in the revealed preferences from $\mathcal{D}'$ would imply a cycle in $\succ$, a contradiction. Moreover, all the bundles $x \in X$ are unique by construction. Thus the auxiliary data $\mathcal{D}'$ satisfies Matzkin and Richter (1991)'s Theorem 2^∞(a), and there exists a regular (in fact even infinitely differentiable and strictly concave) utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$ that rationalizes $\mathcal{D}'$ in the classic sense. In particular, for any $x \in X$, optimality of the demands requires $MRS^u_{\ell,\ell'}(x) = v_\ell(x)/v_{\ell'}(x)$. Finally, since $v(x) \in C_\varepsilon(x)$, we know by construction that for every $p \in P(x)$,

$$\frac{p_\ell}{p_{\ell'}}(1-\epsilon) \leq MRS^u_{\ell,\ell'}(x) = \frac{v_\ell(x)}{v_{\ell'}(x)} \leq \frac{p_\ell}{p_{\ell'}}\left(\frac{1}{1-\epsilon}\right).$$

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

Theorem 1 shows that testing whether demand data $\mathcal{D}$ is ε-rationalizable amounts to checking whether the collection of restrictions $\mathcal{R}_\varepsilon(\mathcal{D})$ is acyclically satisfiable. But how does one check the latter condition? We now explain how to do this with de Clippel and Rozen (2018)'s procedure for lower-contour set restrictions.

Consider the following procedure, which (if successful) constructs, from the bottom up, the consumer's preference over the set X of demanded bundles. The first step is to check whether there exists a bundle in X, denoted x^1, such that $\Gamma(x^1, v) = \emptyset$ for some $v \in C_\varepsilon(x^1)$. If so, then there is a candidate for worst element in X, and the procedure continues; but if not, then acyclic satisfiability fails. Next, one checks whether there exists a bundle in $X \setminus \{x^1\}$, denoted x^2, for which $\Gamma(x^2, v) \setminus \{x^1\} = \emptyset$ for some $v \in C_\varepsilon(x^2)$. If so, then there is a candidate for the second-worst element in X, and the procedure continues; but if not, then acyclic satisfiability fails. Continuing in this manner, the procedure will enumerate all the elements of X if and only if the set of lower-contour set restrictions $\mathcal{R}_\varepsilon(\mathcal{D})$ is acyclically satisfiable. Moreover, we show that testing takes only polynomially many steps in the number of observations in $\mathcal{D}$.

Theorem 2 *Acyclic satisfiability of $\mathcal{R}_\varepsilon(\mathcal{D})$ is testable in polynomial time in $|\mathcal{D}|$ using de Clippel and Rozen (2018)'s enumeration procedure.*

Proof. de Clippel and Rozen (2018)'s procedure requires one to iteratively find, after already having found a identified a set of $i-1$ bundles $S \subseteq X$ in previous steps, a bundle $x^i \in X \setminus S$ such that $\Gamma(x^i, v) \setminus S = \emptyset$ for some $v \in C_\varepsilon(x^i)$. As there are only $|X| \leq |\mathcal{D}|$ such elements to rank, it suffices to show that in each step, one can determine in polynomial time whether there is a remaining bundle satisfying the selection criterion. Let $\{c^1, \ldots, c^J\} \subseteq \mathbb{R}^L_{++}$ be the generating vectors of $C_\varepsilon(x)$, that is, $C_\varepsilon(x) = \{\sum_{j=1}^J \alpha_j c^j \mid \alpha_j \geq 0 \ \forall j, \ \sum_{j=1}^J \alpha_j > 0\}$. Then the problem in each step is tantamount to checking whether:

(4)

$$\exists\, (\alpha_1, \ldots, \alpha_J) \geq \vec{0} \text{ s.t. } \sum_{j=1}^J \alpha_j > 0 \text{ and } (\sum_{j=1}^J \alpha_j c^j) \cdot (x' - x) > 0, \ \forall\, x' \in X \setminus S, \ x' \neq x.$$

A linear program with weak inequalities is solvable in polynomial time (Karmarkar, 1984) in the size of inputs, $|X|$ (which is at most the number of observations $|\mathcal{D}|$) and J (which is a function of L). The above problem has strict inequalities, but it is equivalent to the following linear programming problem that has only weak

inequalities, and solvable in at most polynomial time in $|\mathcal{D}|$:

$$(5) \qquad \exists\, (\gamma_1, \ldots, \gamma_J) \geq \vec{0} \text{ s.t. } (\sum_{j=1}^{J} \gamma_j c^j) \cdot (x' - x) \geq 1,\ \forall\, x' \in X \setminus S,\ x' \neq x.$$

Solutions to (5) also solve (4). Conversely, if $(\alpha_1, \ldots, \alpha_J)$ solves (4), then $(\frac{\alpha_1}{\sigma}, \ldots, \frac{\alpha_J}{\sigma})$, where $\sigma = (\sum_{j=1}^{J} \alpha_j c^j) \cdot (x' - x)$, solves (5). *Q.E.D.*

Our result shows that testing grows at most polynomially with the size of the data, $|\mathcal{D}|$. As can be seen from the proof, solving the linear program in each step of the enumeration procedure is also polynomial in the number J of generating vectors for the cones $C_\varepsilon(x)$. While it is simple to find a formula to identify the generating vectors, their number grows quickly in L. Of course, the dimensionality of consumption bundles is typically very small in lab experiments. It may be large in real life. However, demand data tends to have different products aggregated into a relatively small number of categories, yielding a relatively low value for L and thus a reasonable value for J. Such aggregation may in fact yield ε-Rationalizability, even if the consumer is rational, since the tradeoffs across artificially-constructed categories may not perfectly reflect the consumer's true tradeoffs.

4. Special Cases

In Sections 4.1-4.3, we consider the testable implications of ε-Rationalizability for three of the most classic, and most often applied, assumptions on the utility function: quasi-linearity, homotheticity, and additive separability.

Recall from Section 3 that acyclic satisfiability is a natural extension of SARP for situations where the revealed restrictions on a preference are more complicated than simple comparisons (e.g. 'x must be preferred to y or z' instead of 'x must be preferred to y'). Since the empirical content of ε-Rationalizability corresponds to such revealed restrictions, acyclic satisfiability and the enumeration procedure provide an intuitive framework for testing. There does not seem to be a natural way to extend these ideas when checking for ε-rationalizability by a utility function that is also, for instance, quasi-linear. Even for standard rationalizability by quasi-linear, homothetic, or additively-separable utility functions, the classic tests are typically described instead by the existence of a solution to certain sets of inequalities. This

makes the tests somewhat harder to interpret, though still practically useful.[8]

This section illustrates how the classic, inequality-based approach extends to ε-Rationalizability for these special classes of utility functions, and how the introduction of ε relaxes the inequalities in each case. Key variables in the inequalities correspond to the gradient of the consumer's utility function at the observed choices. Under rationality, these *unobservable* gradients have to be proportional to the *observable* price vector. This relationship breaks down when introducing ε, but prices nonetheless reveal a range of possible gradients. The graphical approach we take for quasi-linear and homothetic utilities, based on shifting chosen bundles along the income-offer curve so as to fall in the same indifference class, makes the inequalities easier to understand. For additively-separable utilities, we gain intuition by building on Rubinstein and Salant (2012)'s idea of a decision maker who uses only 'nearby' preferences, which in this context sheds light on how to interpret misperception of MRS.

4.1 Quasi-linear Utility

A special case of interest is that of a consumer with quasi-linear preferences. Suppose there is an $(L+1)^{st}$ good, *money*, in which the preference is linear. In this case, the market value of the non-monetary goods consumed need not exhaust the consumer's income. As in most situations, we assume that the modeler does not have access to information about the consumer's monetary endowment, either before or after consumption. The analysis would remain unchanged even if these variables were observed.

The demand data $\mathcal{D}$ is *quasi-linearly rationalizable* if there exists a regular utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$ such that for each $(p, x) \in \mathcal{D}$, there are positive monetary endowments m_B (before consumption) and m_A (after consumption) for which (x, m_A) solves

$$\max_{\{(y,m)\in\mathbb{R}^{L+1}_+ \mid p\cdot y+m\leq p\cdot x+m_B\}} u(y) + m.$$

Following the same motivation as the general case, we say that $\mathcal{D}$ is *quasi-linearly ε-rationalizable* if there exists a regular utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$ such that

$$(6)\quad \frac{p_\ell}{p_{\ell'}}(1-\varepsilon) \leq MRS^u_{\ell\ell'}(x) \leq \frac{p_\ell}{p_{\ell'}}\Big(\frac{1}{1-\varepsilon}\Big) \quad \text{and} \quad p_\ell(1-\varepsilon) \leq \frac{\partial u}{\partial x_\ell}(x) \leq p_\ell\Big(\frac{1}{1-\varepsilon}\Big),$$

[8] Similarly, we could prove Theorem 1 by developing a generalization of Afriat's inequality-based approach for ε-Rationalizablity, but the resulting test would be less intuitive than the one proposed.

for each $\ell \neq \ell'$ in $\{1, \ldots, L\}$ and each observation $(p, x) \in \mathcal{D}$. In this case, u is said to *quasi-linearly ε-rationalize* the demand data. The consumer's utility function over extended bundles (x, m) takes the form $u(x) + m$, and must satisfy the same conditions as before over all $L + 1$ goods. Conditions involving the $(L + 1)^{st}$ good (money) simplify to the second type of inequality above.

To understand the testable implications of quasi-linear ε-rationalizability, it is instructive to start with $\varepsilon = 0$ (Rationality). If $u : \mathbb{R}^L_+ \to \mathbb{R}$ quasi-linearly rationalizes $\mathcal{D}$, then one can augment each observation (p, x) with a monetary level $m(p, x) \geq 0$ so that the consumer becomes indifferent among all augmented bundles:

$$u(x) + m(p, x) = u(x') + m(p', x'),$$

for all (p, x) and (p', x') in $\mathcal{D}$.[9] Importantly, quasi-linearity ensures that MRS's are invariant to monetary levels, as the indifference curves are all parallel shifts along the monetary dimension. Using this fact, strict quasi-concavity of u implies

$$p \cdot x + m(p, x) < p \cdot x' + m(p', x'), \tag{7}$$

for all (p, x) and (p', x') in $\mathcal{D}$ such that $x \neq x'$. Figure 1(a) illustrates this reasoning when $L = 1$.

Differentiability implies a second necessary condition. Optimality of the augmented bundle requires in particular that the bundle of non-monetary goods is preferred over all cheaper alternatives: for all $(p, x) \in \mathcal{D}$, $u(x) \geq u(y)$ for all y such that $p \cdot y \leq p \cdot x$. Differentiability of u implies that $\mathcal{D}$ may contain two observations where the same bundle is consumed at two price vectors, only if the budget sets for non-monetary goods coincide, that is,

$$\text{If } (p, x),\ (p', x) \in \mathcal{D}, \text{ then } p \text{ and } p' \text{ are co-linear.} \tag{8}$$

The following observation, which is proved in the Appendix, establishes that these inequalities are not only necessary but also sufficient.[10] In fact, beyond regularity, one

[9] Indeed, simply fix some $\bar{u}$ larger or equal than the utility of any bundle x appearing in $\mathcal{D}$, and take $m(x, p) = \bar{u} - u(x)$ for all $(x, p) \in \mathcal{D}$.

[10] This observation is similar to Brown and Calsamiglia (2007, Theorem 2.2), with a couple of differences. First, they look for a quasi-linear rationalization with a continuous, concave and strictly monotone utility (as opposed to regularity). As such, they do not have condition (8) and their

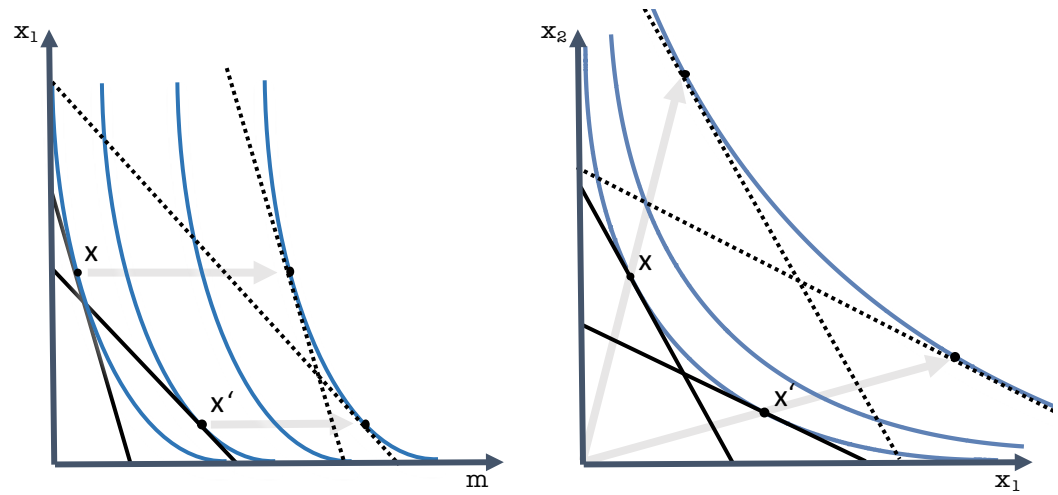

(a) Quasi-linear utility with $L = 1$.

(b) Homothetic utility with $L = 2$.

Figure 1: Illustrations of the construction in the case $\varepsilon = 0$.

can even find a utility function that is strictly concave and infinitely differentiable.

Observation 1 *The demand data $\mathcal{D}$ satisfies (7) for some $m : \mathcal{D} \to \mathbb{R}_+$ and (8) if and only if there exists a regular $u : \mathbb{R}^L_+ \to \mathbb{R}$ that quasi-linearly rationalizes $\mathcal{D}$.*

The above reasoning extends to positive ε's, but we must first identify candidate marginal rates of substitution from the data. Observing $(p, x) \in \mathcal{D}$ informs us that the gradient of the true utility function $u(\cdot)+m$ at x belongs to the convex polyhedron $C^{ql}_\epsilon(p, x)$ defined by the inequalities in (6). Contrary to (8), the same bundle x may be demanded under price vectors generating different budget sets. Such observations tell us that the gradient of the true utility function $u(\cdot)+m$ at x belongs to the convex polyhedron $C^{ql}_\varepsilon(x) = \cap_{\{p|(p,x)\in\mathcal{D}\}} C^{ql}_\varepsilon(p, x)$.

After conjecturing a gradient $\lambda(x) \in C^{ql}_\varepsilon(x)$ for each demanded bundle x, the same reasoning as in the case $\varepsilon = 0$ identifies the following necessary condition: in addition to finding the λ's as above, one must find nonnegative monetary levels $m(p, x)$ for every observation $(p, x) \in \mathcal{D}$, such that these jointly satisfy:

$$\lambda(x) \cdot x + m(p, x) < \lambda(x) \cdot x' + m(p', x'), \tag{9}$$

version of (7) has only weak inequalities. Second, the proofs are different. Only sufficiency remains in our case, which follows fairly simply from Matzkin and Richter (1991, Theorem 2^∞); by contrast, they show necessity using subgradients of concave functions and sufficiency using Rockafeller (1970)'s cyclical monotonicity and his related construction. Finally, our proof shows that strict quasi-concavity and strict concavity cannot be empirically distinguished in this setting.

for all $(p, x), (p', x') \in \mathcal{D}$ for $x \neq x'$. We now show these conditions are also sufficient.

Theorem 3 *The demand data $\mathcal{D}$ is quasi-linearly ε-rationalizable if and only if (9) is satisfied for some choices of $\lambda(x) \in C_\varepsilon^{ql}(x)$ and $m(p, x) \geq 0$ for each $(p, x) \in \mathcal{D}$.*

Proof. It remains to establish sufficiency. Suppose the condition is satisfied, and construct the auxiliary data $\mathcal{D}' = (\lambda(x), x)_{x \in X}$. By construction, the demanded bundles in $\mathcal{D}'$ are all unique, so $\mathcal{D}'$ trivially satisfies condition (8). For each $x \in X$, fix $p(x)$ such that $(p(x), x) \in \mathcal{D}$, and let $m(\lambda(x), x) = m(p(x), x)$. Then $\mathcal{D}'$ satisfies (7), since $\mathcal{D}$ satisfies (9). By Observation 1, $\mathcal{D}'$ is quasi-linearly rationalizable using some regular utility function $u : \mathbb{R}_+^L \to \mathbb{R}$. Quasi-linear ε-rationalizability of $\mathcal{D}$, using this same u, follows from $\lambda(x) \in C_\varepsilon^{ql}(x)$ for all $x \in X$. *Q.E.D.*

Note that testing quasi-linear ε-rationalizability amounts to solving a linear programming problem with some strict inequalities. As before, the problem is equivalent to one with weak inequalities (see Lemma 1 in the Appendix), and thus solvable in polynomial time in the number of observations, $|\mathcal{D}|$.

4.2 Homothetic Utility

Another special case of interest is that of a consumer with homothetic preferences. Say that a utility function u is homothetic if it is homegenous of degree 1: $u(\alpha x) = \alpha u(x)$ for each $\alpha > 0$ and each bundle $x \in \mathbb{R}_+^L$. The demand data $\mathcal{D}$ is *homothetically ε-rationalizable* if Definition 1 applies with the added restriction that u is homothetic.

It is instructive once again to first understand the testable implications with $\varepsilon = 0$ (Rationality). If $u : \mathbb{R}_+^L \to \mathbb{R}$ homothetically rationalizes $\mathcal{D}$, then one can rescale each demanded bundle x by a factor $\alpha(x) \geq 1$ so that the consumer becomes indifferent among all rescaled bundles:

$$\alpha(x)u(x) = \alpha(x')u(x'),$$

for all $x, x' \in X$.[11] Importantly, homotheticity ensures that MRS's are invariant to rescaling. Using this fact, strict quasi-concavity of u implies

$$\alpha(x)p \cdot x < \alpha(x')p \cdot x', \tag{10}$$

[11] Indeed, simply fix some $\bar{u}$ larger or equal than the utility of any bundle $x \in X$, and take $\alpha(y) = \bar{u}/u(y)$ for all $y \in X$ ($u(y) > 0$, since homotheticity and strong monotonicity imply that u is strictly positive on $\mathbb{R}_{++}^L$).

for all (p, x) and (p', x') in $\mathcal{D}$ such that $x \neq x'$. Figure 1(b) illustrates this reasoning when $L = 2$. Differentiability also implies (8), just as in the quasi-linear case.

The following observation, which is proved in the Appendix, shows that these inequalities are also sufficient.[12] In fact, beyond regularity, one can even find a utility function that is strictly concave and infinitely differentiable.

Observation 2 *The demand data $\mathcal{D}$ satisfies (8) and (10), for some $\alpha : X \to [1, \infty)$, if and only if there is a regular $u : \mathbb{R}^L_+ \to \mathbb{R}$ that homothetically rationalizes $\mathcal{D}$.*

The above ideas extends to positive ε's, but we must first use the data to identify candidate marginal rates of substitution. As a demanded bundle $x \in X$ might be chosen at different price vectors, the data informs us that the gradient of the true utility function u at x belongs to the cone $C_\varepsilon(x)$ defined in (2). After finding a gradient $\lambda(x)$ for each $x \in X$, the reasoning from the case $\varepsilon = 0$ may be repeated to yield the following necessary condition: in addition to finding the λ's as above, one must find $\alpha(x) \geq 1$ for all $x \in X$, such that these jointly satisfy:

$$\alpha(x)\lambda(x) \cdot x < \alpha(x')\lambda(x) \cdot x', \tag{11}$$

for all $x, x' \in X$ with $x \neq x'$. We now show these conditions are also sufficient. Thus, testing homothetic ε-rationalizability amounts to a *bilinear* programming problem, which in this case is solvable in polynomial time (see Lemma 2 in the Appendix).

Theorem 4 *The demand data $\mathcal{D}$ is homothetically ε-rationalizable if and only if (11) is satisfied for some choices of $\lambda(x) \in C_\varepsilon(x)$ and $\alpha(x) \geq 1$ for each $x \in X$.*

Proof. It remains to establish sufficiency. Suppose the condition is satisfied, and construct the auxiliary data $\mathcal{D}' = (\lambda(x), x)_{x \in X}$. By construction, the demanded bundles in $\mathcal{D}'$ are all unique, so $\mathcal{D}'$ trivially satisfies condition (8). Moreover, $\mathcal{D}'$ satisfies (10) because the α's and λ's satisfy (11). By Observation 2, $\mathcal{D}'$ is homothetically rationalizable by a regular utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$. Homothetic ε-rationalizability of $\mathcal{D}$, using this same u, follows from $\lambda(x) \in C_\varepsilon(x)$ for all $x \in X$. *Q.E.D.*

[12] Observation 2 is related to Varian (1983, Theorem 2), as the inequalities (10) are similar to those in the third item of his theorem (Varian normalizes prices by total expenditure, making this similarity less obvious). The most relevant difference is that he does not require *strict* concavity/quasi-concavity, which leads to weak inequalities and no condition (8). Consequently, our proof is different from Varian's, instead mimicking our previous argument for quasi-linearity. Again, it follows fairly simply from Matzkin and Richter (1991, Theorem 2^∞), but in this case uses their construction to obtain only one representative indifference curve, that is used to build a homothetic utility function.

4.3 Additive Separability

Consider a consumer whose utility function is additively separable. Say that a regular utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$ is additively separable if there exist strictly concave,[13] differentiable and strictly monotone utility functions $u_\ell : \mathbb{R}_+ \to \mathbb{R}$ for each ℓ, such that $u(x) = \sum_{\ell=1}^L u_\ell(x_\ell)$ for all $x \in \mathbb{R}^L_+$. Simple addition of the u_ℓ's is without loss of generality, as they can absorb scaling factors. The demand data $\mathcal{D}$ is *additively ε-rationalizable* if Definition 1 applies with the restriction that u is additively separable.

In this context, we develop an interesting observation relating to Rubinstein and Salant (2012)'s notion of a decision maker who uses only preferences that are 'close' to her true one. In particular, additive ε-rationalizability is equivalent to the consumer making each choice optimally using a utility function which has an analogous additive structure: it uses the same u_ℓ's as her true utility function u, but coefficients scaling the u_ℓ's twist her MRS to a limited extent.

Observation 3 *$\mathcal{D}$ is additively ε-rationalizable by $u(\cdot) = \sum_{\ell=1}^L u_\ell(\cdot)$ if and only if there exists $\alpha : \mathcal{D} \to \mathcal{R}_{++}$ such that for all $(p, x) \in \mathcal{D}$,*

$$x = \arg\max_{\{p\cdot y \le p\cdot x\}} \sum_{\ell=1}^{L} \alpha_\ell(p,x) u_\ell(x_\ell), \text{ and} \tag{12a}$$

$$1-\varepsilon \le \frac{\alpha_\ell(p,x)}{\alpha_{\ell'}(p,x)} \le \frac{1}{1-\varepsilon}, \text{ for all } \ell, \ell' \in \{1, \ldots, L\}. \tag{12b}$$

Proof. We begin with necessity. Given regularity, note that condition (12a) holds for $x \in \mathbb{R}^L_{++}$ if and only if the first-order conditions for optimality hold. Thus optimality for the additive utility function in (12a) holds by setting $\alpha_\ell(p,x)/\alpha_{\ell'}(p,x) := (p_\ell/p_{\ell'}) MRS^u_{\ell',\ell}(x)$. The bounds (12b) are implied by ε-Rationalizability using u.

For sufficiency, (12a) implies that for each $(p, x) \in \mathcal{D}$, the optimality condition $(\alpha_\ell(p,x)/\alpha_{\ell'}(p,x)) MRS^u_{\ell,\ell'}(x) = p_\ell/p_{\ell'}$ holds for all ℓ, ℓ'. Using (12b), it is easy to see that $u(\cdot) = \sum_{\ell=1}^L u_\ell(\cdot)$ additively ε-rationalizes $\mathcal{D}$. *Q.E.D.*

Consider, for instance, an intertemporal interpretation where x_ℓ is the amount of good ℓ in time-period ℓ, and the consumer's true preferences are given by discounted

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

utility (*i.e.*, $u_\ell = \delta^\ell \tilde{u}$ for a utility function $\tilde{u}$ which is common across all ℓ's). Then Observation 3 means that all the consumer's errors could be attributed to misperception of discounting functions by a scaling factor, in the spirit of quasi-hyperbolic discounting. Similarly, in a setting with risk, where each ℓ is a state of the world, all her errors could be attributed to misperception of probabilities.

Observation 3 is also helpful for understanding the testable implications of additive ε-Rationalizability. Indeed, if the demand data $\mathcal{D}$ is additively ε-rationalizable by u, then Observation 3 allows us to find a weighting function α so that each observation (p, x) is rational for the utility function $\sum_\ell \alpha_\ell(p, x) u_\ell(y_\ell)$. Then, by extending the ideas from Varian (1983) for standard rationalizability by an additively separable utility function, we prove the following characterization in the Appendix. As before, testing the conditions is polynomial in $|\mathcal{D}|$ (see Lemma 3 in the Appendix). Recalling that X is the set of demanded bundles, let $X_\ell = \{x_\ell \mid x \in X\}$.

Theorem 5 *$\mathcal{D}$ is additively ε-rationalizable if and only if there exist $\beta : \mathcal{D} \to \mathbb{R}^L_{++}$ and $z_\ell : X_\ell \to \mathbb{R}$ for each ℓ such that*

$$z_\ell(x'_\ell) < z_\ell(x_\ell) + \beta_\ell(p, x) p_\ell(x'_\ell - x_\ell) \text{ if } x_\ell \neq x'_\ell, \tag{13a}$$

$$\beta_\ell(p, x) p_\ell = \beta_\ell(p', x') p'_\ell \text{ if } x_\ell = x'_\ell, \text{ and} \tag{13b}$$

$$1 - \varepsilon \leq \frac{\beta_\ell(p, x)}{\beta_{\ell'}(p, x)} \leq \frac{1}{1 - \varepsilon}, \tag{13c}$$

for all goods ℓ, ℓ' and all $(p, x), (p', x') \in \mathcal{D}$.

We recover Varian's classic result when $\varepsilon = 0$ (Rationality), in which case $\beta_\ell(p, x)$ is the same across all goods by equation (13c). To gain some intuition, it can be seen from the proof in the Appendix that $\beta_\ell(p, x) p_\ell$ represents the modeler's guess about the consumer's marginal utility of u_ℓ at x_ℓ. Utility maximization means it must be proportional to p_ℓ, with the proportionality constant (the Lagrange multiplier) independent of ℓ. Introducing ε adds a degree of freedom, allowing for variations across goods in the limited range given in (13c). Notice, however, that the modeler's guess cannot contradict itself in the event that the consumer picks the same quantity of good ℓ in two different observations, hence requirement (13b).

A special case of interest, relevant for applications to discounted utility (exponential or not) and expected utility, is when there is a common $\tilde{u}(\cdot) : \mathbb{R}_+ \to \mathbb{R}$ such that $u_\ell(\cdot) = \delta(\ell)\tilde{u}(\cdot)$ for each ℓ. Theorem 5 and its proof can be adapted to obtain testable implications in those cases as well.

5. Measuring Irrationality through Tradeoff Misperception

ε-Rationalizability limits the degree to which marginal rates of substitutions and opportunity costs can differ. It naturally lends itself to a new way of measuring the extent of irrationality in demand data; namely, a *Tradeoff-Misperception Index* may be defined by taking the infimum over all ε such that the data is ε-rationalizable.[14]

This provides an alternative measure to Afriat (1973)'s Critical Cost Inefficiency Index, which computes the minimal factor by which the consumer's budgets must be shrunk to eliminate all revealed-preference cycles. To formalize Afriat's index, define for $\sigma \in [0,1]$ a strict revealed preference $x \succ_{A,\sigma} y$ if $(p,x) \in \mathcal{D}$ and $\sigma p \cdot x > p \cdot y$; and a weak revealed preference $x \succeq_{A,\sigma} y$ if $\mathcal{D}$ contains a sequence of observations $(p^1, x^1), \ldots, (p^n, x^n)$ where $x^1 = x$, $x^n = y$, and for each $i \in \{1, \ldots, n-1\}$, either $x^i = x^{i+1}$ or $\sigma p^i \cdot x^i \geq p^i \cdot x^{i+1}$. Then the Critical Cost Inefficiency Index is the infimum of $1 - \sigma$ over those $\sigma \in [0,1]$ such that $x \succeq_{A,\sigma} y$ implies not $y \succ_{A,\sigma} x$ for all observed choices x, y.[15]

In this section, we develop a deeper understanding of the Tradeoff-Misperception Index. First, we establish a surprising relationship with Afriat's index. Second, we discuss benefits and drawbacks of our index versus money metrics in general. Finally, following Halevy et al (2018), we show how the Tradeoff-Misperception Index can also be used to study the goodness of fit of parametric classes of preferences.

[14] While it would be difficult to find the Tradeoff-Misperception Index with arbitrary precision, it takes only n applications of the enumeration procedure to identify the index with $\pm\frac{1}{2^n}$ precision. First, test for 1/2-rationalizability; then test 1/4-rationalizability if the previous test succeeds, and 3/4-rationalizability otherwise; and continue this recursively $n-2$ more times.

[15] Varian (1990) generalizes Afriat's index to allow the (proportional) budget adjustment to vary per observation. Houtman and Maks (1985) propose another classic index, the smallest subset of the data that needs to be dropped to make it rationalizable. Halevy et al (2018) points out that these indices can be viewed as being in the same umbrella class, with different restrictions on the possible budget adjustments, and different ways of aggregating across observations. We refer the reader to Halevy et al (2018, Appendix B) for a thorough discussion of some other metrics that don't fall under this umbrella, such as Echenique, Lee and Shum (2011) and Apesteguia and Ballester (2015). One can also consider measures of irrationality by creating a demand function through interpolation and using the norm of the resulting Slutsky matrix (Aguiar and Serrano, 2017).

5.1 The Relation to Afriat's Critical Cost Inefficiency Index

The Tradeoff-Misperception Index measures local errors in assessing the marginal rates of substitution. Afriat's Critical Cost Inefficiency Index, by contrast, examines global monetary effects: what percentage of income is lost by making irrational choices? While these measures may appear unrelated, there is a perhaps unexpected connection between the two.

Theorem 6 *For any $\varepsilon \geq 0$ and any demand data $\mathcal{D}$, if the Tradeoff-Misperception Index is less than or equal to ε, then Afriat's Critical Cost Inefficiency Index is also less than or equal to ε. However, the converse is false.*

Theorem 6 establishes that the Tradeoff-Misperception Index is more demanding than the Critical Cost Index: a small error in perception of MRS means only small budgetary adjustments are needed to eliminate revealed preference cycles, but not vice-versa. The first statement in the result is proved below; and a counterexample to its converse is given in Example 1 below that. Using Halevy et al (2018, Theorem 1), a corollary of Theorem 6 is that demand data which is ε-rationalizable will satisfy their notion of $\mathbf{v}$-GARP, where $\mathbf{v} = (1-\varepsilon, \ldots, 1-\varepsilon)$, ensuring that there will be a locally non-satiated (even concave) utility function which rationalizes the revealed preferences remaining after all incomes are scaled by $1-\varepsilon$ à la Afriat to eliminate cycles.

Proof. Theorem 1 shows $\mathcal{D}$ is ε-rationalizable if and only if $\mathcal{R}_\varepsilon(\mathcal{D})$ is acyclically satisfiable. As seen from the proof of that result, acyclic satisfiability of $\mathcal{R}_\varepsilon(\mathcal{D})$ implies that for every $x \in X$, there is $v(x) \in C_\varepsilon(x)$ such that the auxiliary demand data $\mathcal{D}' = (v(x), x)_{x \in X}$ is rationalizable (in the classic sense) by a regular utility function u. We will use u to show that Afriat's index for the original data $\mathcal{D}$ is at most ε.

Take any $(p, x) \in \mathcal{D}$, and consider the indifference curve of u passing through x. As illustrated in Figure 2(a), the budget set may not be tangent to the indifference curve at x, but the line determined by the vector $v(x)$ is tangent, given the rationalizability of $\mathcal{D}'$. We claim that $\frac{p \cdot y}{p \cdot x} \geq 1-\varepsilon$ for any bundle y above this tangent line, i.e., any y such that $v(x) \cdot y \geq v(x) \cdot x$. This holds trivially if $p = v(x)$, so suppose they are different and consider the optimization problem

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

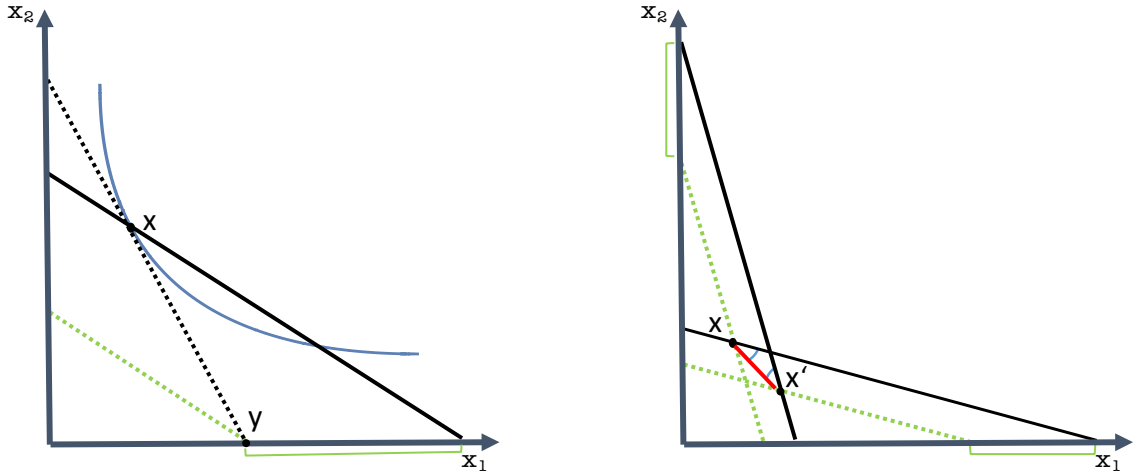

(a) The construction in the proof.

(b) Illustration for Example 1.

Figure 2: Understanding the relationship with Afriat's index in Theorem 6.

The constraint must bind at the optimum, else the objective could be further reduced. Also, as seen in Figure 2(a), linearity of the objective and constraint imply the solution must occur at a bundle y with only one positive component: that is, there is ℓ such that $y_\ell = \frac{v(x)\cdot x}{v_\ell(x)}$ and $y_i = 0$ for all $i \neq \ell$. Using the fact that $v(x) \in C_\varepsilon(x)$, the minimal expenditure satisfies:

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

By quasi-concavity of u, any bundle z with $u(z) \geq u(x)$ must satisfy $v(x)\cdot z \geq v(x)\cdot x$. Hence, the above inequality shows that if each budget set is scaled down by $1-\varepsilon$, the choice x from the original budget set is strictly preferred, according to u, to all bundles in the shrunken budget set. To finish the proof that the Critical Cost Inefficiency Index is at most ε, observe that any cycle in $\succeq_{A,1-\varepsilon}$ would imply a cycle in the corresponding utilities generated by u, which is impossible. *Q.E.D.*

We now show by example that Afriat's index can be strictly smaller than our own; and that, moreover, there exists demand data where the two indices coincide.

Example 1 *Consider demand data with two observations: the bundle $x = (a, b)$ is chosen at the price vector $p = (1, \pi)$, and the bundle $x' = (b, a)$ is chosen at the price vector $p' = (\pi, 1)$, where $\pi > 1$. This situation is illustrated in Figure 2(b). We*

first claim that the Tradeoff Misperception Index is $1-\frac{1}{\pi}$*. To see this, remember from the proofs of Theorems 1 and 6 that* ε*-Rationalizability ensures there is a selection of* $v(\tilde{x}) \in C_\varepsilon(x)$ *for each* $\tilde{x} \in X$ *so that the auxiliary data* $(v(\tilde{x}), \tilde{x})_{\tilde{x}\in X}$ *is rationalizable in the classic sense by a regular utility function* u*. Hence* $v(x)$ *(respectively,* $v(x')$*) determines the slope of the tangent to the indifference curve of* u *at the bundle* x *(respectively,* x'*). The red line in the figure, which connects* x *and* x'*, has slope* -1*. If* $\frac{v_1(x)}{v_2(x)} \leq 1$*, then the tangent to the indifference curve of* u *at the bundle* x *would be flatter than the red line, and strict quasi-concavity of* u *would imply that* $u(x) > u(x')$*. Analogously, if* $\frac{v_1(x')}{v_2(x')} \geq 1$*, then* $u(x') > u(x)$*. Hence at least one of these relationships must be false. A necessary condition for* $\frac{v_1(x)}{v_2(x)} > 1$*, and in fact also for* $\frac{v_1(x')}{v_2(x')} < 1$*, is* $\varepsilon > 1-\frac{1}{\pi}$ *(e.g.,* $1 < \frac{v_1(x)}{v_2(x)} \leq \frac{1}{\pi}\frac{1}{1-\varepsilon}$ *since* $v(x) \in C_\varepsilon(x)$*). Conversely, if* $\varepsilon > 1-\frac{1}{\pi}$ *then we can rationalize the auxiliary data using a regular* u*. To see this, notice that* $MRS^u_{12}(x)$ *must belong to* $(\frac{1-\varepsilon}{\pi}, \frac{1}{1-\varepsilon}\frac{1}{\pi})$ *and* $MRS^u_{12}(x')$ *must belong to* $((1-\varepsilon)\pi, \frac{\pi}{1-\varepsilon})$*. The condition on* ε *amounts to* $(1-\varepsilon)\pi < 1$*, so* $MRS^u_{12}(x)$ *can be larger than* 1 *and* $MRS^u_{12}(x')$ *can be smaller than* 1*. Hence there is room to construct a regular utility function* u *while avoiding a WARP violation, by making the tangent at* x *steeper than, and the tangent at* x' *flatter than, the red line connecting them.*

Noteworthily, the Tradeoff Misperception Index is $1-\frac{1}{\pi}$ *independently of the values* a, b*. For any* x *and* x' *which are oppositely placed at the frontier of their respective budget sets, these choices correspond to the same sizable error in local perception of MRS. By contrast, Afriat's index becomes arbitrarily small as* x, x' *get arbitrarily close to the intersection of the two budget sets. Thus the Tradeoff-Misperception Index is more demanding: data considered near-rational under Afriat's measure might not be considered near-rational under our measure. There also exists demand data where the two measures agree. To see this, shift* x *to the vertical axis and* x' *to the horizontal axis. Since* $p \cdot x = p' \cdot x'$ *and* $x' = (p' \cdot x'/\pi, 0)$*, Afriat's index is given by* $1 - \frac{p \cdot x'}{p \cdot x} = 1 - \frac{p' \cdot x'/\pi}{p \cdot x} = 1 - \frac{1}{\pi}$*.*

5.2 Comparisons of Robustness

The Tradeoff-Misperception Index is based on ε-Rationalizability, which relies on finding a regular utility function. The definition of regularity could be generalized at the cost of more burdensome notation to require only quasi-concavity and strict monotonicity; but ε-Rationalizability loses bite without quasi-concavity, as the definition is based on first-order conditions. In this dimension, Afriat's index and other

money metrics have a benefit over our index, in that they are robust to whether underlying preferences are convex.

However, a benefit of the Tradeoff-Misperception Index over Afriat's index, and similar money metrics, is that our concept is robust to more general settings. Indeed, it is well defined even for nonlinear budget sets and other domains of choice.[16] While the Afriat index is clearly defined for linear budget sets, one can imagine at least three possible ways to extend it to nonlinear ones.

To illustrate, suppose Afriat's index is one-half: does this mean we shrink by 1/2 the budget frontier, the consumer's endowment, or the consumer's monetary income? For a simple example where these differ, suppose the consumer has \$10 to spend on two goods: apples, which cost \$1 each, and bananas, which cost \$1 each for up to five bananas and \$0.50 for each additional one. The quantity discount on bananas means the consumer's budget set is non-linear. Proportionally shrinking the budget set seems intuitive, as it retains the original shape; but this leads to a different outcome than shrinking monetary income. Indeed, with only \$5 to spend, the budget set would simply be linear. Of course, in non-monetary settings, a consumer's income is given by the value of their endowment. In the classic setting with linear prices, one can halve the budget by halving the endowment. Suppose we halve the consumer's endowment in our example. If the consumer is an apple farmer who brings 10 apples to the market, then he faces a linear budget with only 5 apples (corresponding to an 'income' of \$5). But if the consumer is a banana farmer who brings 15 bananas to the market, then she still faces a nonlinear budget set with only 7.5 remaining (corresponding to an 'income' of \$7.50). Thus shrinking the endowment may not lead to a robust solution.

Such concerns would not arise for the Tradeoff-Misperception Index, as it only considers how well an individual understands her marginal-utility tradeoffs, and involves no adjustments to constraint sets. For other domains of choice, such as players deciding what strategy to use in a game, it is unclear how to apply a money metric at all. By contrast, the Tradeoff-Misperception Index applies whenever a regular utility can be defined on the domain of choice.

5.3 Measuring Goodness of Fit

Halevy et al (2018) investigates the use of money metrics for assessing the good-

[16] Rational choice from nonlinear budget sets is studied by Forges and Minelli (2009).

ness of fit of parametric classes of preferences. They define loss functions for a general class of money metrics, and show that these can be additively decomposed into two components: one which measures irrationality and one which captures preference misspecification. As formalized below, analogous results apply for our notion of tradeoff misperception.

For any two positive vectors $x, y \in \mathbb{R}^L_{++}$, define

$$\varepsilon(x,y) := \min\{\varepsilon \mid \varepsilon \in [0,1] \text{ and } 1-\varepsilon \leq \frac{x_\ell/x_{\ell'}}{y_\ell/y_{\ell'}} \leq \frac{1}{1-\varepsilon}, \forall\, \ell, \ell' = 1, \ldots, L\}.$$

To understand goodness of fit for a class of utility function, first consider a single regular utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$. The *Tradeoff-Misperception vector* for a regular utility function u, $\nu_{tm}(\mathcal{D}, u)$ associates to each $(p, x) \in \mathcal{D}$ the number $\varepsilon(p, \nabla u(x))$. Let the *Tradeoff-Misperception Index for* u, $I_{tm}(\mathcal{D}, u)$, be the largest component of the vector $\nu_{tm}(\mathcal{D}, u)$. This index picks up the largest distance (measured in terms of tradeoff misperception) between the price vector and the marginal rate of substitution across all observations.[17] The *Tradeoff-Misperception Index for a class* $\mathcal{U}$ of regular utility functions is $I_{tm}(\mathcal{D}, \mathcal{U}) = \inf\{I_{tm}(\mathcal{D}, u) \mid u \in \mathcal{U}\}$.

Let $\mathcal{U}^*$ denote the set of all regular utility functions, and let $I^*_{tm}(\mathcal{D})$ be a shortcut for $I_{tm}(\mathcal{D}, \mathcal{U}^*)$. Then we have the following observation.

Observation 4 *The following statements hold:*

(i) Take any $u \in \mathcal{U}^*$ and $\varepsilon \in [0, 1)$. Then u ε-rationalizes the consumer data $\mathcal{D}$ if and only if $I_{tm}(\mathcal{D}, u) \leq \varepsilon$.

(ii) If the consumer data $\mathcal{D}$ is ε-rationalizable then $I^*_{tm}(\mathcal{D}) \leq \varepsilon$. Conversely, if $I^*_{tm}(\mathcal{D}) \leq \varepsilon$ then $\mathcal{D}$ is ε'-rationalizable for all $\varepsilon' > \varepsilon$.

(iii) If $\mathcal{U} \subseteq \mathcal{U}'$, then $I_{tm}(\mathcal{D}, \mathcal{U}') \leq I_{tm}(\mathcal{D}, \mathcal{U})$. In particular, $I^*_{tm}(\mathcal{D}) \leq I_{tm}(\mathcal{D}, \mathcal{U})$.

(iv) $I_{tm}(\mathcal{D}, \mathcal{U}) = I^*_{tm}(\mathcal{D}) + [I_{tm}(\mathcal{D}, \mathcal{U}) - I^*_{tm}(\mathcal{D})]$, for any $\mathcal{U} \subseteq \mathcal{U}^*$.

Parts (i) and (ii) clearly follow from Definition 1, and formalize a straightforward connection between the Tradeoff-Misperception index and ε-Rationalizability. Part

[17] One can think of many other ways to aggregate the vector into an index, as in Halevy et al. (2018) who introduces an aggregation function f. For concreteness and notational simplicity, we restrict attention to the max aggregator. A similar worst-case scenario approach singles out the Afriat inconsistency index among Varian's class of money metric indices.

(ii) takes into account that the class of regular utility functions is open due to strict quasi-concavity. These results lead to a simple decomposition in part (iv) of I_{tm} for any class $\mathcal{U}$ of regular utility functions in terms of (1) tradeoff misperception with respect to the whole class of regular utility functions, and (2) any remaining discrepancy due to misspecification arising from the use of a smaller class $\mathcal{U}$. Part (iii), which ensures that the discrepancy from misspecification is nonnegative, follows from monotonicity of the infimum in the class of utility functions.

6. Misperception in Prices

In our model, consumers perceive prices accurately, but may misperceive utility tradeoffs at the margin. Market prices inform the modeler of the subjective tradeoff the consumer used to reach her decision. This provides bounds on the MRS of the true underlying preference, which can be used to formulate testable implications. Recent extensions of consumer theory have instead explored the implications of standard utility maximization when prices are misperceived. This section draws connections between this approach and ours, and shows how our testing methodology can prove useful in the case of misperceived prices as well.

Let u be a consumer's utility function, let p be the true price vector, let p^s be the price vector perceived by the consumer, and let e be her endowment. Presumably, the modeler observes p and perhaps e, but neither u nor p^s. Testable implications arise after positing how p^s crystallizes as a function of p. For the moment, let us discuss consumer choices while treating p^s as a known exogenous variable. Things seem simple at first: the consumer picks the u-maximal bundle given what she thinks she can afford ($p^s \cdot x \leq p^s \cdot e$). Bounded rationality in this sense simply boils down to rational choice applied to misperceived prices. Unfortunately, this approach is ill-defined, as the real budget constraint involves p and not p^s. When it comes to paying for the chosen bundle, there may be too little money to pay for it, or some income may be left unused.

As Gabaix (2014) puts it, we would naturally like to think of an "agent [who] is boundedly rational, but smart enough to exhaust his budget." One can imagine different adjustment methods to make the consumer's choice fall on the actual (p) budget line; see Chetty, Looney and Kroft (2007) and Gabaix (2014).[18] Gabaix

[18]The 2007 working-paper version of Chetty et al. (2009) has a more complete discussion of this topic than the published version.

(2014) explores the case of a consumer who exhausts her true income and uses perceived prices when comparing opportunity costs to marginal rates of substitution. As illustrated in Figure 3(a), choices are obtained by identifying where the income-offer curve for p^s crosses the actual budget line.[19] We refer to this as the *first adjustment rule*. Another option, mentioned in Chetty et al. (2007) and Gabaix (2014), is to adjust the consumption level of a pre-specified good so as to shift the consumer's chosen bundle from the p^s to the p budget line. This is illustrated in Figure 3(b). We refer to this as the *second adjustment rule*. Choosing which good to adjust is somewhat ad hoc in general, with the exception of quasi-linear settings, where the numeraire is a natural choice. In that case, however, the two adjustment rules coincide.

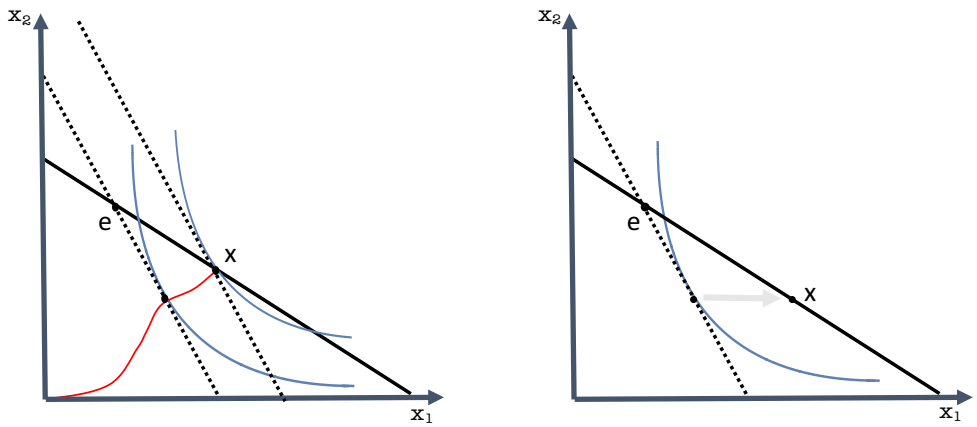

(a) The first adjustment rule.

(b) The second adjustment rule.

Figure 3: Illustrations of two possible adjustment rules for regaining budget balance under the correct prices. The dotted lines correspond to misperceived price ratios.

We now return to the question of how p^s crystallizes. Gabaix's (2014) sparsemax model captures a decision maker who, for instance, realizes that spending time understanding the state of the Amazonian forest or interest rates in some distant country will be costly and yet have very little effect on her decision. The decision maker knows default values of such variables (e.g., long-run averages) and optimally allocates effort in determining what price estimate, somewhere between the default value and the true value, to obtain. Formally, denoting the default price vector by p^d, the perceived price for good ℓ is $p^s_\ell = p^d_\ell + m_\ell(p_\ell - p^d_\ell)$, where $m_\ell \in [0, 1]$ represents the

[19] While consumer theory constitutes an important application, it is worth noting that Gabaix develops his theory for even more general decision problems.

extent to which the consumer shifts the perceived price of the good from the default value to the actual price. While Gabaix's theory endogenously determines the vector $m = (m_1, \ldots, m_L)$,[20] we suggest a variant where any $m \in [0,1]^L$ is conceivable:

$$p^s \in C(p, p^d) = \{(p^d_\ell + m_\ell[p_\ell - p^d_\ell])_{\ell=1}^L \mid m \in [0,1]^L\}. \tag{14}$$

Clearly, the perceived prices generated by Gabaix's theory belong to $C(p, p^d)$. Thus, requiring price vectors to belong to $C(p, p^d)$ leads to predictions that are less precise than Gabaix's, but robust to a wide variety of theories of subjective price formation.

Suppose that the modeler observes the consumer picking x under the price vector p, while conjecturing that default prices are given by p^d. Under the first adjustment rule discussed above, this observation reveals that the gradient of u at x is proportional to some $p^s \in C(p, p^d)$. This defines a different convex set of possible gradients than for ε-Rationalizability, with the key exogenous variables being p and p^d instead of p and ε. The same testing methodology applies, with analogous lower contour-set restrictions, simply by replacing $C_\varepsilon(p, x)$ by the new set of possible gradients. While Gabaix (2014) derives properties of the Slutsky matrix of the full demand function associated to the sparse-max model, this adaptation of Theorem 1 provides testable implications in the spirit of Afriat for a more permissive theory of subjective prices.

It is also interesting to consider a consumer whose utility is quasi-linear in the $(L+1)^{st}$ good, given our observation that the two adjustment rules coincide in this case. Without loss of generality, the price of good $L+1$ is normalized to 1 (and never misperceived), in which case it is understood that the price vectors p, p^d and p^s for the other goods are expressed in terms of this numeraire. We will focus once again on the case where p^s belongs to $C(p, p^d)$. When the modeler witnesses the consumer picking the bundle $(x, m) \in \mathbb{R}^{L+1}$, the consumer's expected consumption level in good $L+1$ cannot be known for sure: it is equal to $p^s \cdot (e - x) + e_{L+1}$, but the modeler does not know p^s. As should be clear from the analysis in Section 4.1, quasi-linearity makes this lack of information irrelevant, as the test we provide for quasi-linearity does not require the modeler to have access to information about the consumer's monetary endowment, either before or after consumption. Then, as for the first adjustment rule, testing can be performed simply by replacing $C^{ql}_\varepsilon(p, x)$ by

[20] This corresponds to Step 1 of his theory. By contrast, the analysis of the previous paragraph, as summarized in Figure 3(a), corresponds to Step 2.

the appropriate set of possible gradients derived from $C(p, p^d)$ as 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.

Aguiar, V. and R. Serrano (2017), Slutsky Matrix Norms: The Size, Classification, and Comparative Statics of Bounded Rationality, *Journal of Economic Theory*, 172, 163–201.

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

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

Cherepanov, V., T. Feddersen and A. Sandroni (2013), Rationalization, *Theoretical Economics*, 8, 775–800.

Chetty, R., A. Looney and K. Kroft (2007), Salience and Taxation: Theory and Evidence, *NBER Working Paper no. 13330*.

Chetty, R., A. Looney and K. Kroft (2009), Salience and Taxation: Theory and Evidence, *American Economic Review*, 99, 1145–1177.

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

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

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

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

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

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

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

Halevy, Y., D. Persitz and L. Zrill (2018), Parametric Recoverability of Preferences, *Journal of Political Economy, forthcoming*.

Kalai, G., A. Rubinstein, and R. Spiegler (2002), Rationalizing Choice Functions by Multiple Rationales, *Econometrica*, 70(6), 2481–2488.

Lleras, J., Y. Masatlioglu, D. Nakajima, and E. Ozbay (2017), When More is Less: Choice by Limited Consideration, *Journal of Economic Theory*, 170, 70–85.

Manzini, P., and M. Mariotti (2007), Sequentially Rationalizable Choice, *American Economic Review*, 97(5), 1824–1839.

Manzini, P., and M. Mariotti (2012), Categorize Then Choose: Boundedly Rational Choice and Welfare, *Journal of the European Economics Association*, 10(5), 1141–1165.

Masatlioglu, Y., D. Nakajima and E. Ozbay (2012), Revealed Attention, *American Economic Review*, 102(5), 2183–2205.

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

Rockafeller, T. (1970), *Convex Analysis*, Princeton University Press.

Rubinstein, A. and Y. Salant (2012), Eliciting Welfare Preferences from Behavioral Data Sets, *Review of Economic Studies*, 79, 375–387.

Steiner, J. and C. Stewart (2016), Perceiving Prospects Properly, *American Economic Review*, 106, 1601–1631.

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

Varian, H.R. (1983), Non-Parametric Tests of Consumer Behaviour, *The Review of Economic Studies*, 50(1), 99–110.

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

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

Appendix

Proof of Observation 1 (Testable implications of quasi-linear rationalizability)

Necessity was established in the text. As for sufficiency, observe that the assumptions of Theorem 2^∞(c) of Matzkin and Richter (1991) hold with $\lambda^i = 1$, thanks to (7) and (8). Hence, there exists a regular utility function u (in fact, even strictly concave and infinitely differentiable) $u : \mathbb{R}^L_+ \to \mathbb{R}$ such that the gradient of u at x is p, for each $(p, x) \in \mathcal{D}$.[21] Thus u quasi-linearly rationalizes $\mathcal{D}$, as desired. *Q.E.D.*

Lemma 1 *Testing if (9) is satisfied for some choices of $\lambda(x) \in C^{ql}_\varepsilon(x)$ and $m(p, x) \geq 0$ for each $(p, x) \in \mathcal{D}$ is equivalent to a linear program with weak inequalities.*

Proof. First observe that for each x, $\lambda(x) \in C^{ql}_\varepsilon(x)$ if and only if $\lambda(x) = \sum_j \alpha_j(x)c^j(x)$ for some $\alpha(x) = (\alpha_1(x), \ldots, \alpha_J(x)) \geq 0$ with $\alpha(x) \cdot \vec{1} = 1$, where the $c^j(x)$'s are the extreme points of $C^{ql}_\varepsilon(x)$. We claim the original problem, with λ's replaced using α's and c^j's as above, is equivalent to finding $\gamma : X \to \mathbb{R}^J_+$ and $n : \mathcal{D} \to \mathbb{R}_+$ such that $\gamma(x) \cdot \vec{1} = \gamma(x') \cdot \vec{1}$ and $(\sum_j \gamma_j(x)c^j(x)) \cdot (x' - x) + (n(p', x') - n(p, x)) \geq 1$, for all $(p, x), (p', x')$ in $\mathcal{D}$ with $x \neq x'$. Given a solution to the original problem, define $\sigma = \min_{(p,x),(p',x') \in \mathcal{D}, x' \neq x} (\sum_j \alpha_j(x)c^j(x)) \cdot (x' - x) + (m(p', x') - m(p, x)) > 0$. Then γ, n given by $\gamma(x) = \alpha(x)/\sigma$ and $n(p, x) = m(p, x)/\sigma$ for all (p, x) solve the second problem. Now suppose γ, n solve the second problem. Observe that it is impossible to have $\gamma(x) \cdot \vec{1} = 0$ for some x (and thus all x); otherwise $\gamma_j(x) = 0$ for all j and x, implying $n(p', x') - n(p, x) \geq 1$ for all $x \neq x'$, a contradiction. As $\gamma(x) \cdot \vec{1} > 0$ for all $x \in X$, define $\alpha(x) = \gamma(x)/(\gamma(x) \cdot \vec{1})$ for all $x \in X$. Then α, n solve the original problem, proving the desired equivalence. *Q.E.D.*

Proof of Observation 2 (Testable implications of homothetic rationalizability)

Necessity was established above. As for sufficiency, observe that the assumptions of Theorem 2^∞(c) of Matzkin and Richter (1991) hold for the observations $\{\alpha(x)x | x \in X\}$ with (in their notations) $u^i = 0$ and $\lambda_i = 1$, thanks to (8) and (10). Hence,

[21] To see this, first look at the proof of Lemma 2 in Matzkin and Richter's paper. Their equation (4.19a) tells that the gradient of ϕ_i at a demanded bundle x^i is $\lambda^i p^i$. Then inequality (4.15a) tells that the minimum defining U in equation (4.18) is achieved by ϕ_i in a neighborhood of x^i. Finally, using (5.1c) in the proof of Theorem 2^∞, a convolution argument à la Chiappori and Rochet (1987) provides an infinitely differentiable utility function u that preserves these properties. Given our inequalities (7), $\lambda^i = 1$ for all data points i, and the gradient of u at x is p for all $(p, x) \in \mathcal{D}$.

there exists a regular utility function $w : \mathbb{R}^L_+ \to \mathbb{R}$ (in fact, even strictly concave and infinitely differentiable) such that $w(\alpha(x)x) = u^i = 0$ and the gradient of w at $\alpha(x)x$ is p, for each $x \in X$.[22] Let IC_0 be the zero-utility indifference curve, that is, all bundles $y \in \mathbb{R}^L_+$ such that $w(y) = 0$. For all $z \in \mathbb{R}^L_+ \setminus \{\vec{0}\}$, let $u(z) = w(\beta z)/\beta$, where $\beta > 0$ is the unique scalar such that $\beta z \in IC_0$.[23] Let $u(\vec{0}) = 0$. Clearly, u inherits regularity from w, and homothetically rationalizes $\mathcal{D}$. *Q.E.D.*

Lemma 2 *Testing if (11) is satisfied for some choices of $\lambda(x) \in C_\varepsilon(x)$ and $\alpha(x) \geq 1$ for each $x \in X$ is equivalent to a tractable bilinear program.*

Proof. First, note this problem is equivalent to the following one with weak inequalities: find $\beta(x) \geq 1$ and $\mu(x) \in C_\varepsilon(x)$ for each $x \in X$ such that $1 + \beta(x)\mu(x) \cdot x \leq \beta(x')\mu(x) \cdot x'$ for all $x, x' \in X$ with $x \neq x'$. A solution to this new problem clearly solves (11). Conversely, take α's and λ's solving (11); then for each $x \neq x'$, $\alpha(x)\lambda(x) \cdot x + \delta(x, x') = \alpha(x')\lambda(x) \cdot x'$, where $\delta(x, x') > 0$. Dividing the equality for each x by $\min_{x' \in X \setminus \{x\}} \delta(x, x') > 0$, and setting $\beta(x) = \alpha(x)$ and $\mu(x) = \lambda(x)/\min_{x' \in X \setminus \{x\}} \delta(x, x')$ (which must still belong to the cone $C_\varepsilon(x)$), we obtain a solution to the new problem. Finally, using the same trick as Brown and Calsamiglia (2007) when they study (standard) rationalizability by a utility that is *both* quasilinear and homothetic, the change of variables $\beta(x) = \exp(z(x))$ results in smooth, convex inequalities that should be solvable in polynomial time in $|\mathcal{D}|$ by interior point methods. *Q.E.D.*

Proof of Theorem 5 (Testable implications of additive ε-rationalizability)

We begin with necessity. Suppose $\mathcal{D}$ is additively ε-rationalizable by $u(\cdot) = \sum_{\ell=1}^L u_\ell(\cdot)$. For each observation (p, x), let $\alpha(p, x) \in \mathbb{R}^L_{++}$ be a vector of weights as in Observation 3. Since x maximizes $\sum_\ell \alpha_\ell(p, x) u_\ell(y_\ell)$ for all $y \in \mathbb{R}^L_+$ such that $p \cdot y \leq p \cdot x$, there exists a Lagrange multiplier $\gamma(p, x) > 0$ (independent of ℓ) such that

$$\frac{du_\ell}{dx_\ell}(x_\ell) = \frac{\gamma(p, x) p_\ell}{\alpha_\ell(p, x)}, \text{ for all goods } \ell. \tag{15}$$

[22] The explanation in Footnote 21 applies here too.

[23] To check that it is well-defined, notice first that $w(\vec{0}) < 0$ since w is strictly monotone and there are larger bundles with zero utility. Also, one can check that $w(\beta z)$ tends to infinity as β tends to infinity. To see this, note that the partial derivative of w with respect to each good is larger than a strictly positive constant, thanks to (4.19b) in Matzkin and Richter (1991) and the properties of the convolution they use to prove their Theorem 2^∞. For the interested reader, that convolution argument is available in greater detail in Lemma 2^∞ of the 1987 working paper. By the intermediate value theorem, there exists β such that $\beta z \in IC_0$; and strict monotonicity of w ensures uniqueness.

For each good ℓ, by strict concavity of u_ℓ, it must be that

$$u_\ell(y_\ell) < u_\ell(x_\ell) + \frac{du_\ell}{dx_\ell}(x_\ell)(y_\ell - x_\ell), \text{ for all } y \in \mathbb{R}^L_+ \text{ with } y_\ell \neq x_\ell.$$

By (15), this is equivalent to

$$u_\ell(y_\ell) < u_\ell(x_\ell) + \frac{\gamma(p,x)}{\alpha_\ell(p,x)} p_\ell(y_\ell - x_\ell), \text{ for all } y \in \mathbb{R}^L_+ \text{ with } y_\ell \neq x_\ell.$$

In particular, this inequality must apply to any $y \in X$ with $y_\ell \neq x_\ell$. The necessary conditions (13a) and (13c) follow by taking $z_\ell(x_\ell) = u_\ell(x_\ell)$, for each $x \in X$, and $\beta_\ell(p,x) = \gamma(p,x)/\alpha_\ell(p,x)$, for each $(p,x) \in \mathcal{D}$. Finally, (13b) follows from equation (15), as $\frac{\gamma(p,x)p_\ell}{\alpha_\ell(p,x)}$ must equal $\frac{\gamma(p',x')p'_\ell}{\alpha_\ell(p',x')}$ if $x_\ell = x'_\ell$.

We now show sufficiency. Given the inequalities (13a), there exists for each good ℓ a strictly concave, infinitely differentiable and strictly monotone utility function $u_\ell : \mathbb{R}_+ \to \mathbb{R}$ such that $u_\ell(x_\ell) = z_\ell(x_\ell)$ and $du_\ell(x_\ell)/dx_\ell = \beta_\ell(p,x)p_\ell$ for each $(p,x) \in \mathcal{D}$ (which is well-defined by (13b)). Writing $\alpha_\ell(p,x)$ for $1/\beta_\ell(p,x)$, this implies that for each $(p,x) \in \mathcal{D}$, the bundle x maximizes $\sum_{\ell=1}^L \alpha_\ell(p,x)u_\ell(y_\ell)$ for all $y \in \mathbb{R}^L_+$ such that $p \cdot y \leq p \cdot x$. Clearly, the α's inherit property (13c) from the β's, and we conclude from Observation 3 that $\mathcal{D}$ is additively ε-rationalizable by $u(\cdot) = \sum_{k=1}^K u_\ell(\cdot)$. *Q.E.D.*

Lemma 3 *Testing if (13a)–(13c) are satisfied for some choices of $\beta : \mathcal{D} \to \mathbb{R}^L_+$ and $z_\ell : X_\ell \to \mathbb{R}$ for each ℓ is equivalent to a linear program with weak inequalities.*

Proof. We prove that the problem is equivalent to the following linear program with weak inequalities: find $\hat{\beta} : \mathcal{D} \to \mathbb{R}^L_+$ and $\hat{z}_\ell : X_\ell \to \mathbb{R}$ for each ℓ such that $1 \leq \hat{z}_\ell(x_\ell) - \hat{z}_\ell(x'_\ell) + \hat{\beta}_\ell(p,x)p_\ell(x'_\ell - x_\ell)$ if $x_\ell \neq x'_{\ell'}$, $\hat{\beta}_\ell(p,x)p_\ell = p'_\ell\hat{\beta}_\ell(p',x')$ if $x_\ell = x'_{\ell'}$ and $1 - \varepsilon \leq \frac{\hat{\beta}_\ell(p,x)}{\hat{\beta}_{\ell'}(p,x)} \leq \frac{1}{1-\varepsilon}$ hold for all $(p,x),(p',x') \in \mathcal{D}$ and all ℓ. A solution to this new problem clearly solves the original, since $\hat{\beta}$ must be strictly positive to satisfy the ratio inequalities. Conversely, take β and z_ℓ's satisfying the original problem. Since there are finitely many inequalities in (13a) and they are all strict, one can find $\eta > 0$ such that they continue to hold when adding η to the left-hand side. Setting $\hat{\beta}_\ell(p,x) = \beta_\ell(p,x)/\eta$ and each $\hat{z}_\ell(x_\ell) = z_\ell(x_\ell)/\eta$ solves the new problem. *Q.E.D.*