

Revealed Price Preference: Theory and Empirical Analysis

RAHUL DEB

University of Toronto

YUICHI KITAMURA

Yale University

JOHN K.-H. QUAH

Johns Hopkins University and National University of Singapore

and

JÖRG STOYE

Cornell University

First version received October 2018; Editorial decision December 2021; Accepted July 2022 (Eds.)

To determine the welfare implications of price changes in demand data, we introduce a revealed preference relation over prices. We show that the absence of cycles in this relation characterizes a consumer who trades off the utility of consumption against the disutility of expenditure. Our model can be applied whenever a consumer's demand over a strict subset of all available goods is being analysed; it can also be extended to settings with discrete goods and non-linear prices. To illustrate its use, we apply our model to a single-agent data set and to a data set with repeated cross-sections. We develop a novel test of linear hypotheses on partially identified parameters to estimate the proportion of the population who are revealed better off due to a price change in the latter application. This new technique can be used for non-parametric counterfactual analysis more broadly.

Key words: Revealed preference, Welfare analysis, Non-parametric counterfactual analysis.

JEL Codes: C14, D11, D12, D60

1. INTRODUCTION

A central issue in economic analysis is the determination of the welfare effect of price changes. As an example, suppose we observe a consumer's purchases of two goods, gasoline and food, from two separate trips to a grocery store with an on-site gasoline retailer. In the first instance t , the prices are $p^t = (2, 2)$ of gasoline and food, respectively and she buys a bundle $x^t = (6, 3)$. In her second trip t' , the prices are $p^{t'} = (3, 1)$ and she purchases $x^{t'} = (5, 4)$. The most basic welfare question one can ask here is whether the consumer is better off at the prices prevailing at t or at t' (keeping fixed the prices of all other goods she consumes)? In this article, we introduce a

The editor in charge of this paper was Aureo de Paula.

theoretical framework based on revealed preference, along with a non-parametric econometric technique, that would allow us to answer questions of this type.

A typical approach to this problem is to model the consumer as having a quasilinear utility $\tilde{U}(x) - p \cdot x$ since, in particular, this allows for simple “sufficient statistics” analysis of welfare gains or losses using a Harberger formula (see Chetty, 2009 and most recently Kleven, 2021 for an overview of this approach). The second term ($-p \cdot x$) in the quasilinear utility captures the fact that the goods being analysed (food and gasoline in our example) do *not* constitute the universe of the consumer’s consumption, so that expenditure lowers utility because it reduces the consumption of an outside (numeraire) good.

The point of departure of our analysis is the simple observation that, even without modelling the consumer’s preference as quasilinear (or taking any other functional form) we can still conclude that she is better off at t compared to t' . This is because $p^{t'} \cdot x^{t'} = 19$ whereas $p^t \cdot x^{t'} = 18$. If the prices at t' were p^t instead of $p^{t'}$, the consumer would be better off since purchasing the same bundle $x^{t'}$ would cost less, leaving her with more money to buy other goods (apart from gasoline and food).¹ More generally, the consumer has a *preference over prices* that an analyst could at least partially discern from the data: if at observations t and t' , we find that $p^t \cdot x^{t'} \leq (<) p^{t'} \cdot x^{t'}$, then

the consumer has revealed that she (strictly) prefers the price p^t to the price $p^{t'}$.

Welfare comparisons made in this way are consistent only if the revealed preference relation over prices has no cycles, a property we call the *generalized axiom of price preference* (GAPP). A natural question then arises: what does GAPP mean for consumer behaviour?

1.1. Augmented utility functions

To answer this question, we assume that the analyst collects a data set $\mathcal{D} = \{p^t, x^t\}_{t=1}^T$ from a consumer; each observation t consists of the prices $p^t \in \mathbb{R}_{++}^L$ of L goods (representing some but not all the goods she consumes) and the consumer’s demand $x^t \in \mathbb{R}_+^L$ at those prices. We show that GAPP (on \mathcal{D}) is both necessary and sufficient for the existence of a strictly increasing function $U: \mathbb{R}_+^L \times \mathbb{R}_- \rightarrow \mathbb{R}$ that *rationalizes* \mathcal{D} in the following sense:

$$x^t \in \operatorname{argmax}_{x \in \mathbb{R}_+^L} U(x, -p^t \cdot x) \text{ for all } t = 1, 2, \dots, T.$$

The function U should be interpreted as an *expenditure-augmented utility function*, where $U(x, -e)$ is the consumer’s utility when she purchases x after spending e . Note that the consumer’s optimal expenditure on the observed goods is dependent on prices: she could in principle spend more than what she actually spent (as she optimizes over $x \in \mathbb{R}_+^L$) but the trade-off is the disutility of greater expenditure. Observe that quasilinear utility $U(x, -p \cdot x) = \tilde{U}(x) - p \cdot x$ is a special case of an augmented utility function.

Below are features of the augmented utility model that make it widely applicable.

(1) Being more general, it does not have the overly strong implications of the quasilinear model on consumer demand (see Section 2.3). In particular, it is broad enough to accommodate several

1. Another way of seeing this is the following. Suppose t' is a supermarket where the prices are $p^{t'}$ and we observe the bundle $x^{t'}$ being bought by a consumer. If at supermarket t , the prices are p^t , then we know that the consumer would prefer this supermarket, since the same purchases at t' would cost less at t .

prominent behavioural economics models such as reference dependence, mental budgeting, and inattention to prices. We briefly describe the first of these here, Section 2.4 has a more detailed discussion. Kőszegi and Rabin (2006) and Heidhues and Kőszegi (2008) argue that consumption decisions can depend not just on the actual prices but also on the prices the consumer *expected* to pay. Specifically, the disutility from spending is greater if the expected price was lower than the sticker price and vice versa. A simple way they propose of capturing this phenomenon is the following function

$$U(x, -px) = \tilde{U}(x) - p \cdot x - F(p \cdot x - \tilde{p} \cdot x).$$

The first two terms capture standard quasilinear preferences whereas the third term captures a general form of reference dependence.² In Koszegi and Rabin's terminology, the consumer gets "gain-loss utility" by comparing the expenditure $p \cdot x$ she incurs on a bundle x against the expenditure $\tilde{p} \cdot x$ she expected to incur, where \tilde{p} are her reference prices.³ While the introduction of the third term means that this utility is no longer quasilinear, it remains an augmented utility function (when \tilde{p} is kept fixed).

(2) Since a consumer's utility at prices p is given by $\max_{x \in \mathbb{R}_+^L} U(x, -p \cdot x)$, this obviously leads to a ranking or preference on prices. Going further, it is possible to develop notions analogous to compensating and equivalent variations which gives us a *quantitative* sense of how much one set of prices is preferred to another (see Section 3.3).

(3) Readers familiar with Afriat's Theorem (Afriat, 1967) will no doubt have already noticed that we are working in a similar framework. That theorem characterizes a data set $\mathcal{D} = \{p^t, x^t\}_{t=1}^T$ that could be rationalized in the following sense: there is $\tilde{U}: \mathbb{R}_+^L \rightarrow \mathbb{R}$ such that $\tilde{U}(x^t) \geq \tilde{U}(x)$ for all $x \in \mathbb{R}_+^L$ that satisfy $p^t \cdot x \leq p^t \cdot x^t$. The notion of rationalization in our model is distinct from that in Afriat's Theorem (even the utility functions have different domains) and there are data sets that could be rationalized in one sense but not the other. We explain these differences in Sections 2.2 and 3.1.

(4) In Section 4, we show that our framework generalizes to accommodate discrete choice, characteristics models, and non-linear prices.

1.2. Random augmented utility model

In the second part of the article, we develop the random version of the augmented utility model, in order to study the demand distribution of a population of consumers drawn from repeated cross-sectional data. We first devise a test to check if the data are consistent with the random augmented utility model (RAUM). We then develop a procedure to *estimate* the *proportion* of consumers who are made better or worse off by a given change in prices; welfare analysis of this kind under general preference heterogeneity is a challenging empirical issue, and has attracted considerable recent research (see e.g. Hausman and Newey, 2016 and its references).

Unlike the case of data collected from a single individual, it is worth noting that, in this case, both model testing and welfare analysis are statistical since we need to account for sampling error inherent in repeated cross sectional data. Our RAUM test uses existing (though recently developed) econometric methods. On the other hand, to carry out the welfare analysis, we have to develop new theoretical econometric results. This is a standalone contribution that has applications beyond this article.

2. For a related model of reference prices leading to a similar functional form, see Sakovics (2011).

3. A common choice for F is $F(p \cdot x - \tilde{p} \cdot x) = \max\{\bar{k}(p \cdot x - \tilde{p} \cdot x), 0\} + \min\{\underline{k}(p \cdot x - \tilde{p} \cdot x), 0\}$ with $\bar{k} > \underline{k} > 0$ or that the consumer feels losses relative to the reference point more severely than commensurate gains.

We argue that testing the RAUM on actual repeated cross-sectional data (such as household survey data) is easier than testing the random utility model (RUM), that is, the random version of the standard model where consumers maximize utility subject to a budget constraint. A test for the RUM is described in McFadden and Richter (1991) (henceforth MR), but two hurdles must be overcome before their test can be implemented. First, MR do not account for finite sample issues as they assume that population distributions of demand are observed; recently, Kitamura and Stoye (2018) (henceforth KS) developed a testing procedure which incorporates sampling error. Second, the test suggested by MR requires the observation of large samples of consumers who face not only the same prices but also make *identical* total expenditures (on the observed goods). This feature is not true of any real observational data where a consumer's demand (and thus total expenditure) on the observed goods is typically price dependent. Thus to implement their test, KS first estimate demand distributions at a fixed (median) level of total expenditure, which requires the use of an instrumental variable technique (with all its attendant assumptions) to adjust for the dependence of the total expenditure on prices.

In contrast, the RAUM can be tested directly on household survey data, even when the demand distribution at a given price vector implies *heterogeneous levels of total expenditure across consumers*.⁴ This allows us to estimate the demand distribution by simply using sample frequencies, and we can avoid the above-mentioned additional layer of demand estimation needed for testing the RUM.

The reason for this remarkable simplification is somewhat ironic: we show that a data set is consistent with the RAUM if, and only if, a converted version of the data set which has identical total expenditures among consumers at each price passes the RUM test devised by MR. In other words, we apply the test suggested by MR, but not for the model they have in mind. This trick also means that we can use, and in a more straightforward way, the econometric techniques in KS.

Further, we can evaluate the welfare impact of an observed prices change. If we observe the true distribution of demand at each price, theoretical bounds can be derived for the population proportion who are revealed better or worse off after a price change. Of course, for finite samples, these bounds instead have to be estimated. To do so, we develop new econometric techniques to estimate confidence intervals on these population proportions; our methodology builds on the econometric theory in KS but is novel.

We emphasize that these new econometric techniques can be more generally applied to linear hypothesis testing of parameter vectors that are partially identified, even in models that are unrelated to demand theory (see e.g. Lazzati, Quah and Shirai, 2018). They provide a new method for estimation and inference in non-parametric counterfactual analysis and, since the evaluation of counterfactuals is an important goal of empirical research, they are potentially very useful to practitioners.

1.3. Empirical applications

We use separate data sets to demonstrate how welfare analysis can be done using both the deterministic and random versions of our model. First, we use the deterministic model to analyse panel data from the Mexican conditional cash transfer program Progresa. We show that price changes by sellers in response to the cash transfers benefit the untreated households. Second, we conduct welfare analysis with the RAUM using repeated cross-sectional data on household

4. We allow two demands x, x' in the support of the demand distribution at prices p^l to satisfy $p^l \cdot x \neq p^l \cdot x'$.

expenditures from Canada and the UK. This demonstrates how to operationalize our novel econometric methodology to conduct inference for counterfactuals.

2. THE DETERMINISTIC MODEL

We consider an econometrician who is studying a consumer’s demand for L goods. We assume an idealized environment suitable for partial equilibrium analysis, where the consumer’s demand for these goods at different prices are observed, while the consumer’s wealth and the prices of all other goods are held fixed.⁵

Specifically, the econometrician collects a data set with a finite number of observations; each observation t can be represented as (p^t, x^t) , where $p^t \in \mathbb{R}_{++}^L$ are the prices of the L goods and $x^t \in \mathbb{R}_+^L$ is the bundle of those goods purchased by the consumer.⁶ We denote the data set by $\mathcal{D} := \{(p^t, x^t)\}_{t=1}^T$. (We shall slightly abuse notation and use T to refer both to the (finite) number of observations and to the set $\{1, \dots, T\}$; similarly, L could denote both the number, and the set, of commodities.)

We begin with a basic question: given \mathcal{D} , can the econometrician sign the welfare impact of a price change from p^t to $p^{t'}$? Perhaps the most intuitive welfare comparison that can be made is as follows: if at prices $p^{t'}$, the econometrician finds that $p^{t'} \cdot x^t < p^t \cdot x^t$ then he may conclude that the agent is better off at the price vector $p^{t'}$ compared to p^t . This is because, at the price $p^{t'}$ the consumer can, if she wishes, buy the bundle bought at p^t and she would still have money left over to buy other things, so she must be strictly better off at $p^{t'}$. This ranking is eminently sensible, but can it lead to inconsistencies?

Example 1. Consider a two observation data set

$$p^t = (2, 1), x^t = (4, 0) \text{ and } p^{t'} = (1, 2), x^{t'} = (0, 1).$$

Since $p^{t'} \cdot x^t < p^t \cdot x^t$, it seems that the consumer is better off at prices $p^{t'}$ than at p^t ; however, it is also true that $p^t \cdot x^{t'} > p^{t'} \cdot x^{t'}$, which gives the opposite conclusion.

This example shows that a consistent welfare comparison at different prices requires the imposition of some restriction on the data set. To be precise, define the binary relations \succeq_p and \succ_p on $\mathcal{P} := \{p^t\}_{t \in T}$, the set of price vectors observed in \mathcal{D} , in the following manner:

$$p^{t'} \succeq_p (\succ_p) p^t \text{ if } p^{t'} \cdot x^t \leq (<) p^t \cdot x^t.$$

We say that price $p^{t'}$ is *directly (strictly) revealed preferred* to p^t if $p^{t'} \succeq_p (\succ_p) p^t$, that is, whenever the bundle x^t is (strictly) cheaper at prices $p^{t'}$ than at prices p^t . We denote the transitive closure of \succeq_p by \succeq_p^* , that is, for $p^{t'}$ and p^t in \mathcal{P} , we have $p^{t'} \succeq_p^* p^t$ if there are t_1, t_2, \dots, t_N in T such that $p^{t'} \succeq_p p^{t_1}, p^{t_1} \succeq_p p^{t_2}, \dots, p^{t_{N-1}} \succeq_p p^{t_N}$, and $p^{t_N} \succeq_p p^t$; in this case, we say that $p^{t'}$ is *revealed preferred* to p^t . If anywhere along this sequence, \succeq_p can be replaced with \succ_p then we say that $p^{t'}$ is *revealed strictly preferred* to p^t and denote that relation by $p^{t'} \succ_p^* p^t$.⁷ The following restriction,

5. Under fairly standard (but strong) assumptions, changes to the external environment can be precisely justified by deflating the prices of the L goods (see Section 3.5).

6. We postpone the discussion of discrete consumption spaces and non-linear pricing to Section 4.

7. Notice that it makes sense to write $\hat{p} \succeq_p p^t$ even if \hat{p} is not in \mathcal{P} , since the demand at \hat{p} is not needed in the definition revealed preference. Similarly, it is possible to define $\hat{p} \succ_p p^t$ and the transitive extensions $\hat{p} \succeq_p^* p^t$ and $\hat{p} \succ_p^* p^t$. This observation is useful later on, in Sections 3.3 and 5.3.

which excludes circularity in the assessment of the consumer's wellbeing at different prices, is a bare minimum condition to impose on \mathcal{D} .

Definition 2.1. *The data set $\mathcal{D} = \{(p^t, x^t)\}_{t=1}^T$ satisfies the Generalized Axiom of Price Preference or GAPP if there are no observations $t, t' \in T$ such that $p^{t'} \succeq_p^* p^t$ and $p^t \succ_p^* p^{t'}$.*

This in turn leads naturally to the following question: if a consumer's observed demand behaviour obeys GAPP, what could we say about her decision-making procedure?

2.1. The expenditure-augmented utility model

An *expenditure-augmented utility function* or simply, an *augmented utility function*, is a function $U: \mathbb{R}_+^L \times \mathbb{R}_- \rightarrow \mathbb{R}$, where $U(x, -e)$ is the consumer's utility when she spends e to purchase bundle x . We require that $U(x, -e)$ is strictly increasing in the last argument (in other words, is strictly decreasing in expenditure), which captures the tradeoff the consumer faces between consuming x and consuming other goods (outside the set L).

At a given price p , the consumer chooses a bundle x to maximize $U(x, -p \cdot x)$. We denote the *indirect utility at price p* by

$$V(p) := \sup_{x \in \mathbb{R}_+^L} U(x, -p \cdot x). \quad (1)$$

If the consumer's augmented utility maximization problem has a solution at every price vector $p \in \mathbb{R}_{++}^L$, then V is also defined at those prices and this induces a reflexive, transitive, and complete preference over prices in \mathbb{R}_{++}^L .

A data set $\mathcal{D} = \{(p^t, x^t)\}_{t=1}^T$ is *rationalized by an augmented utility function* if there exists such a function $U: \mathbb{R}_+^L \times \mathbb{R}_- \rightarrow \mathbb{R}$ with

$$x^t \in \operatorname{argmax}_{x \in \mathbb{R}_+^L} U(x, -p^t \cdot x) \quad \text{for all } t \in T. \quad (2)$$

It is straightforward to see that GAPP is necessary for a data set to be rationalized by an augmented utility function. First, notice that if $p^{t'} \succeq_p p^t$, then $p^{t'} \cdot x^t \leq p^t \cdot x^t$, and so

$$V(p^{t'}) \geq U(x^t, -p^{t'} \cdot x^t) \geq U(x^t, -p^t \cdot x^t) = V(p^t).$$

Furthermore, $U(x^t, -p^{t'} \cdot x^t) > U(x^t, -p^t \cdot x^t)$ if $p^{t'} \succ_p p^t$, and in that case $V(p^{t'}) > V(p^t)$. Suppose GAPP is not satisfied and there are two observations $t, t' \in T$ such that $p^{t'} \succeq_p^* p^t$ and $p^t \succ_p^* p^{t'}$. Then, there exist $t_1, t_2, \dots, t_N \in T$ which yield the contradiction

$$V(p^{t'}) \geq V(p^{t_1}) \geq \dots \geq V(p^{t_N}) \geq V(p^t) > V(p^{t'}).$$

Our main theoretical result, stated next, also establishes the sufficiency of GAPP for rationalization. Moreover, whenever \mathcal{D} can be rationalized, it can be rationalized by an augmented utility function U with a list of properties that make it convenient for analysis.

Theorem 1. *Given a data set $\mathcal{D} = \{(p^t, x^t)\}_{t=1}^T$, the following are equivalent:*

- (1) \mathcal{D} is rationalized by an augmented utility function.
- (2) \mathcal{D} satisfies GAPP.
- (3) \mathcal{D} is rationalized by an augmented utility function U that is strictly increasing, continuous, and concave. Moreover, U is such that $\max_{x \in \mathbb{R}_+^L} U(x, -p \cdot x)$ has a solution for all $p \in \mathbb{R}_{++}^L$.

2.2. Afriat's Theorem and Proof of Theorem 1

Before we prove Theorem 1, it is worth providing a short description of the standard revealed preference theory of consumer demand and its central result, Afriat's Theorem. This will be useful not just because we will invoke the result several times but also since it will serve as an important point of contrast for our axiom and results.

The standard theory due to Afriat (1967) is built formally on the same primitives as our model: a finite data set of prices and corresponding consumption bundles. Unlike our model however, it is assumed that the observed goods correspond to the *universe* of the consumer's consumption. Formally, a data set \mathcal{D} is said to be *rationalized by a utility function* if there exists a locally non-satiated⁸ utility function $\tilde{U}: \mathbb{R}_+^L \rightarrow \mathbb{R}$ such that

$$x^t \in \operatorname{argmax}_{\{x \in \mathbb{R}_+^L: p^t \cdot x \leq p^t \cdot x^t\}} \tilde{U}(x) \quad \text{for all } t \in T. \quad (3)$$

In words, this criterion asks whether there is a utility function defined over the L observed goods such that the consumer is utility maximizing at every observation t over the fixed budget $p^t \cdot x^t$ corresponding to the observed expenditure.

Of course, data sets (outside of laboratory data) almost never contain the universe of consumed goods and the consumer's *true* budget set is not observed, especially when one takes into account the possibility of borrowing and saving. Given this, when checking if a data set \mathcal{D} can be rationalized in the sense of (3), we are effectively testing whether the consumer is maximizing a sub-utility function $\tilde{U}: \mathbb{R}_+^L \rightarrow \mathbb{R}$ defined specifically on those L goods (or equivalently, has weakly separable preferences).

It should be clear that rationalization in the sense of (3) is distinct from rationalization by an augmented utility function. The augmented utility model specifically takes into account the impact of the prices of these L goods on the consumption of other goods: it is *necessarily* a partial equilibrium model, and designed for partial equilibrium welfare analysis of the type carried out in empirical industrial organization or public economics. An example is the study of the welfare impact of a sales tax levied on a subset of goods.

Revealed preference in Afriat's setting is captured by two binary relations, \succeq_x and \succ_x which are defined as follows on the set $\mathcal{X} := \{x^t\}_{t \in T}$ of chosen bundles observed in \mathcal{D} :

$$x^{t'} \succeq_x (\succ_x) x^t \text{ if } p^{t'} \cdot x^{t'} \geq (>) p^{t'} \cdot x^t.$$

We say that the bundle $x^{t'}$ is *directly revealed (strictly) preferred* to the bundle x^t if $x^{t'} \succeq_x (\succ_x) x^t$, that is, whenever x^t is (strictly) cheaper at prices $p^{t'}$ than $x^{t'}$. This terminology is intuitive: if the agent is maximizing some locally non-satiated utility function $\tilde{U}: \mathbb{R}_+^L \rightarrow \mathbb{R}$, then $x^{t'} \succeq_x x^t$ ($x^{t'} \succ_x x^t$) must imply that $\tilde{U}(x^{t'}) \geq (>) \tilde{U}(x^t)$.

We denote the transitive closure of \succeq_x by \succeq_x^* , that is, for $x^{t'}, x^t \in \mathcal{X}$, we have $x^{t'} \succeq_x^* x^t$ if there are t_1, t_2, \dots, t_N in T such that $x^{t'} \succeq_x x^{t_1}$, $x^{t_1} \succeq_x x^{t_2}$, \dots , $x^{t_{N-1}} \succeq_x x^{t_N} \succeq_x x^t$, and $x^{t_N} \succeq_x x^t$; in this case, we say that $x^{t'}$ is *revealed preferred* to x^t . If any \succeq_x in this sequence can be replaced with \succ_x , we say that $x^{t'}$ is *revealed strictly preferred* to x^t and denote that relation by $x^{t'} \succ_x^* x^t$. Clearly, if \mathcal{D} is rationalizable by some locally non-satiated utility function \tilde{U} , then $x^{t'} \succeq_x^* (\succ_x^*) x^t$ implies that $\tilde{U}(x^{t'}) \geq (>) \tilde{U}(x^t)$. Thus, a necessary condition for \mathcal{D} 's rationalizability is that the revealed preference relation has no cycles.

8. This means that at any bundle x and open neighbourhood of x , there is a bundle x' in the neighbourhood with strictly higher utility.

Definition 2.2. A data set $\mathcal{D} = \{(p^t, x^t)\}_{t=1}^T$ satisfies the Generalized Axiom of Revealed Preference or GARP if there are no observations $t, t' \in T$ such that $x^{t'} \succeq_x^* x^t$ and $x^t \succ_x^* x^{t'}$.

The main claim of Afriat's Theorem is that this condition is also *sufficient* (the formal statement can be found in the Supplementary Appendix A.1.1). Having described Afriat's Theorem, we are now in a position to prove Theorem 1.

PROOF OF THEOREM 1. We will show that (2) \implies (3). We have already argued that (1) \implies (2) and (3) \implies (1) by definition.

Choose a number $M > \max_t p^t \cdot x^t$ and define the augmented data set $\tilde{\mathcal{D}} = \{(p^t, 1), (x^t, M - p^t \cdot x^t)\}_{t=1}^T$. This data set augments \mathcal{D} since we have introduced an $L+1$ th good, which we have priced at 1 across all observations, with the demand for this good equal to $M - p^t \cdot x^t$.

The crucial observation to make here is that

$$(p^t, 1)(x^t, M - p^t \cdot x^t) \geq (p^{t'}, 1)(x^{t'}, M - p^{t'} \cdot x^{t'}) \text{ if and only if } p^{t'} \cdot x^{t'} \geq p^t \cdot x^{t'},$$

which means that

$$(x^t, M - p^t \cdot x^t) \succeq_x (x^{t'}, M - p^{t'} \cdot x^{t'}) \text{ if and only if } p^t \succeq_p p^{t'}.$$

Similarly,

$$(p^t, 1)(x^t, M - p^t \cdot x^t) > (p^{t'}, 1)(x^{t'}, M - p^{t'} \cdot x^{t'}) \text{ if and only if } p^{t'} \cdot x^{t'} > p^t \cdot x^{t'},$$

and so

$$(x^t, M - p^t \cdot x^t) \succ_x (x^{t'}, M - p^{t'} \cdot x^{t'}) \text{ if and only if } p^t \succ_p p^{t'}.$$

Consequently, \mathcal{D} satisfies GAPP if and only if $\tilde{\mathcal{D}}$ satisfies GARP. Applying Afriat's Theorem when $\tilde{\mathcal{D}}$ satisfies GARP, there is $\tilde{U}: \mathbb{R}^{L+1} \rightarrow \mathbb{R}$ (notice that \tilde{U} is defined on \mathbb{R}^{L+1} and not just \mathbb{R}_+^{L+1} ; see Remark 3 in Appendix A.1.1) such that

$$(x^t, M - p^t \cdot x^t) \in \operatorname{argmax}_{\{(x,m) \in \mathbb{R}_+^L \times \mathbb{R} : p^t \cdot x + m \leq M\}} \tilde{U}(x, m) \quad \text{for all } t \in T. \quad (4)$$

The function \tilde{U} can be chosen to be strictly increasing, continuous, and concave, and the lower envelope of a finite set of affine functions. Clearly, the augmented utility function $\tilde{U}: \mathbb{R}_+^L \times \mathbb{R}_- \rightarrow \mathbb{R}$ defined by $\bar{U}(x, -e) := \tilde{U}(x, M - e)$ is strictly increasing in $(x, -e)$, continuous, concave and rationalizes \mathcal{D} .

Define $\hat{U}: \mathbb{R}_+^L \times \mathbb{R}_- \rightarrow \mathbb{R}$ by

$$\hat{U}(x, -e) := \bar{U}(x, -e) - h(\max\{0, e - M\}), \quad (5)$$

where $h: \mathbb{R}_+ \rightarrow \mathbb{R}$ is a differentiable function satisfying $h(0) = 0$, $h'(k) > 0$, $h''(k) \geq 0$ for $k \in \mathbb{R}_+$, and $\lim_{k \rightarrow \infty} h'(k) = \infty$. (For example, $h(k) = k^3$.) Like \bar{U} , the function \hat{U} is strictly increasing in $(x, -e)$, continuous and concave and x^t solves $\max_{x \in \mathbb{R}_+^L} \hat{U}(x, -p^t \cdot x)$ (because $\hat{U}(x, -e) \leq \bar{U}(x, -e)$ for all $(x, -e)$, and $\hat{U}(x^t, -p^t \cdot x^t) = \bar{U}(x^t, -p^t \cdot x^t)$). Furthermore, for every $p \in \mathbb{R}_+^{L+1}$, $\operatorname{argmax}_{x \in X} \hat{U}(x, -p \cdot x)$ is non-empty.⁹ ■

9. Choose a sequence $x^n \in \mathbb{R}_+^L$ such that $\hat{U}(x^n, -p \cdot x^n)$ tends to $\sup_{x \in \mathbb{R}_+^L} \hat{U}(x, -p \cdot x)$ (which we allow to be infinity). It is impossible for $p \cdot x^n \rightarrow \infty$ because the piecewise linearity of $\bar{U}(x, -e)$ in x and the assumption that $\lim_{k \rightarrow \infty} h'(k) \rightarrow \infty$ implies that $\hat{U}(x^n, -p \cdot x^n) \rightarrow -\infty$. So the sequence $p \cdot x^n$ is bounded, which in turn means that there is a subsequence of x^n that converges to $x^* \in \mathbb{R}_+^L$. By the continuity of \hat{U} , we obtain $\hat{U}(x^*, -p \cdot x^*) = \sup_{x \in \mathbb{R}_+^L} \hat{U}(x, -p \cdot x)$.




FIGURE 1

GARP consistent choices that don't allow for consistent welfare predictions.

Note that GARP imposes testable restrictions distinct from GAPP (we postpone a detailed discussion to Section 3.1). This can be seen from Figure 1 which plots not just the observed consumption bundles from Example 1 but also the corresponding budget sets (derived from the observed prices and expenditures). As we argued, GAPP does not hold but, since the budget sets do not even cross, it is clear that GARP does.

From this point onwards, when we refer to “rationalization” without additional qualifiers, we shall mean rationalization by an augmented utility function, that is, in the sense given by (2) rather than in the sense given by (3).

2.3. “Standard” consumer theory and the augmented utility function

Perhaps the clearest motivation for our model is to think of it as a generalization of the quasilinear utility model, in which the consumer derives utility $\tilde{U}(x)$ from the bundle x and maximizes utility net of expenditure, that is, she chooses x to maximize

$$U(x, -e) := \tilde{U}(x) - e, \quad (6)$$

where $e = p \cdot x$. The familiar textbook way of justifying this objective function is to think of the consumer as having a utility function \bar{U} defined over $L + 1$ goods, with the last “outside” good entering additively and linearly into the utility function, so that $\bar{U}(x, z) = \tilde{U}(x) + z$. If the consumer has total wealth W , the utility of buying bundle $x \in \mathbb{R}_+^L$ is then

$$\bar{U}(x, W - p \cdot x) = \tilde{U}(x) - p \cdot x + W.$$

Ignoring boundary issues, the consumer is effectively maximizing (6).

Despite widespread use in partial equilibrium analysis, the complete absence of income effects makes the quasilinear model unsuitable for certain empirical applications. For this reason, it is common to relax the linearity of \bar{U} while retaining the assumption that outside consumption is captured by a *single* outside good; this is true, for example, in the literature on modelling the demand for differentiated goods.¹⁰ Then, the utility of purchasing bundle $x \in \mathbb{R}_+^L$ is $\bar{U}(x, W - p \cdot x)$ and provided W is fixed, the consumer effectively maximizes an augmented utility function: simply let $U(x, -e) = \bar{U}(x, W - e)$.

Obviously, a consumer's outside consumption opportunities would in reality involve more than one good, and the prices of those outside goods could change as well. Within the familiar constrained-optimal model of consumer theory, there are known conditions that justify the representation of those consumption opportunities by a representative good (with its corresponding price index). This is explained in detail in Section 3.5.

Finally, it is worth mentioning that the augmented utility function captures, as a special case, quasilinear utility maximization subject to certain constraints. One example is consumption with a subsistence constraint, which we describe in the empirical application in Section 7.1. Loosely speaking, we can capture constraints on $(x, -e)$ with an augmented utility $U(x, -e)$ that assigns very low values at $(x, -e)$ that violate the constraint.

2.4. Behavioural preferences captured by the augmented-utility model

The central feature of the augmented utility model is that consumers experience disutility from expenditure. As we explained in the previous subsection, this disutility could be interpreted in a purely opportunity cost sense—more expenditure on the consumed goods imply less money available for other goods. In this understanding, the augmented utility function is a reduced form of a broader “true” utility function defined on all goods.

However, it is also reasonable to think that the consumer has—*directly*—a preference over bundles of the observed L goods and their associated expenditure, which she has developed as a way of guiding purchasing decisions. Thus, it is the basic object of analysis and not the reduced form of something more fundamental. This understanding of choice behaviour is exploited in the behavioural economics literature and the following quote from Prelec and Loewenstein (1998) is effectively a description of augmented utility:

each time a consumer engages in an episode of consumption, we assume she asks herself: “How much is this pleasure costing me?” The answer to this question is the imputed cost of consumption. This imputed cost is “real” in the sense that it actually detracts from consumption pleasure.

In this understanding, the disutility of expenditure is still related to opportunity cost, but the relationship is more flexible than what is permitted in a classical framework.

In Section 1, we described how reference-dependent preferences could be captured by an augmented utility function. In the remainder of this section, we describe how our model relates to two other prominent themes in the behavioural literature.

10. For example, in Berry, Levinsohn and Pakes (1995) and in Nevo (2000), \bar{U} is additively separable between the L goods and the outside good; in the former, the utility of consuming y units of the outside good is $\alpha \ln y$, for some $\alpha > 0$, whereas in the latter it is αy (in other words, \bar{U} is quasilinear). In Bhattacharya (2015), \bar{U} is allowed to be a general function defined on $L+1$ goods.

2.4.1. Inattention to prices and expenditure. Chetty, Looney and Kroft (2009) (and the literature that followed) observe that consumers often misperceive prices: specifically, they show that shoppers at grocery stores often do not internalize the price effect of taxes. Gabaix (2019) summarizes this literature and argues that many behavioural biases take the form of inattention.

Our model captures a version of price inattention discussed in Bordalo, Gennaioli and Shleifer (2013) and Gabaix (2014). Here, a consumer perceives the expenditure of a bundle x at prices p as $f(x, p \cdot x)$, where f is increasing in the true expenditure $p \cdot x$ and can potentially depend on x . With this misperception, a consumer with a quasilinear preference chooses $x \in \mathbb{R}_+^L$ to maximize

$$U(x, -p \cdot x) = \tilde{U}(x) - f(x, p \cdot x). \quad (7)$$

A special case is where a consumer has a default price p^d and misperceives the actual price p to be $ap + (1-a)p^d$ where $a \in [0, 1]$ is the “attention parameter”. The perceived expenditure is then $f(x, p \cdot x) = ap \cdot x + (1-a)p^d \cdot x$. More generally, the model accommodates $f(x, p \cdot x) = a(x)p \cdot x + (1-a(x))p^d \cdot x$, where the attention parameter $a(x) \in [0, 1]$ varies across bundles.¹¹ Among other things, this allows a consumer to be more attentive if she is purchasing large bundles compared to small ones (so that $a(x)$ tends to 1 when x is large). Another possibility is that the consumer pays attention to marginal increases in expenditure only when certain thresholds are crossed; this would correspond to the case where f is a step function that depends only on the expenditure $e = p \cdot x$. Clearly, inattention as modelled by (7) is an instance where the agent has an augmented utility function, even though it will typically not be quasilinear (in actual expenditure).¹²

Notice that using an augmented utility function (such as (7)) to capture price inattention is particularly apt because, as noted by Gabaix (2014), the numeraire serves as “the shock absorber that adjusts to the budget constraint”. The alternative is to model the consumer as having *both* price misperception over a given set of goods *and* a budget on those goods that must be satisfied, which inevitably leads to the added complication of modelling how the agent adjusts her intended demand when she realizes it violates (because prices are misperceived) the budget constraint at the true prices.

2.4.2. Budgeting and mental budgeting. As explained in Section 2.3, it is common in partial equilibrium analysis to introduce a numeraire good and assume that the agent has a (standard) utility function and budget set defined on $L+1$ goods, with price and income information used to determine the level of the numeraire consumed. Obviously, this approach requires income information which is not always in the data¹³ and even when it is available, it is strictly speaking not the right value to use as the global budget if the consumer can save and borrow (as acknowledged e.g. in Hausman and Newey (2016)). More generally, determining “the real budget” is not always straightforward, even in a classical setting.

Regularities highlighted by behavioural economists add a further wrinkle to the concept of a budget. It has been widely observed that households do not always treat money as fungible and instead create separate accounts for various categories of goods Thaler (1999).

11. This formulation of perceived expenditure is more general than Gabaix (2014) in that it allows the attention parameter a to depend on x but is less general in that the parameter does not vary across goods.

12. We should add that formulae such as (7) would typically have observable implications that are stronger than GAPP. In other words, price inattention models *do not* have precisely the same empirical content as the augmented utility model.

13. Several widely used data sets, such as supermarket scanner panel data, that contain rich purchase information, do not have accurate income measures. Here, income information is typically the self-reported category (income ranges) when households apply for loyalty cards (so even this information becomes dated).

This is not only true for consumption decisions (see, for instance, Hastings and Shapiro, 2013, 2018) but also for savings decisions, which is why consumers often save more when they have access to commitment savings options (important theoretical and empirical contributions are Amador, Werning and Angeletos (2006) and Feldman (2010), Dupas and Robinson (2013), respectively).

Now consider a researcher modelling the demand for a subset of L goods. If mental accounting effects are important, she must allow for the fact that she cannot observe how the agent categorizes goods, nor does she know the true mental budget that determines expenditure (on the L observed goods and their perceived alternatives). In this situation, augmented utility is a natural way to model the demand for those L goods: it is consistent with constrained utility maximization incorporating an outside good (see Section 2.3) but *does not require* the researcher to take a stand on the unobserved mental budget.¹⁴

3. PROPERTIES OF THE AUGMENTED UTILITY MODEL

In this section, we explore various aspects of the augmented utility model, beginning with a discussion of the relationship between GAPP and GARP.

3.1. Comparing GAPP and GARP

Recall that Example 1 in Section 2 is an example of a data set that obeys GARP but fails GAPP. We now present an example of a data set that satisfies GAPP but fails GARP.

Example 2. Consider the data set consisting of the following two choices:

$$p^t = (2, 1), x^t = (2, 1) \text{ and } p^{t'} = (1, 4), x^{t'} = (0, 2).$$

These choices, as shown in Figure 2, violate GARP as $p^t \cdot x^t = 5 > 2 = p^t \cdot x^{t'}$ ($x^t \succ_x x^{t'}$) and $p^{t'} \cdot x^t = 8 > 6 = p^{t'} \cdot x^{t'}$ ($x^t \succ_x x^{t'}$). However, these choices satisfy GAPP as $p^t \cdot x^t = 5 > 2 = p^t \cdot x^{t'}$ ($p^t \succ_p p^{t'}$) but $p^t \cdot x^t = 5 \not\geq 6 = p^{t'} \cdot x^t$ ($p^{t'} \not\prec_p p^t$).

GAPP and GARP are distinct conditions, but they coincide in data sets where $p^t \cdot x^t = 1$ for all $t \in T$. This is because $x^t \succeq_x (>_x)x^{t'}$ if and only if $p^t \succeq_p (>_p)p^{t'}$ as both require $1 \geq (>)p^t \cdot x^{t'}$. Given a data set $\mathcal{D} = \{(p^t, x^t)\}_{t=1}^T$, we define the *iso-expenditure version* of \mathcal{D} as a data set $\check{\mathcal{D}} := \{(p^t, \check{x}^t)\}_{t=1}^T$, such that $\check{x}^t = x^t / (p^t \cdot x^t)$ and note that $p^t \cdot \check{x}^t = 1$ for all $t \in T$. Observe that the revealed price preference relations $\succeq_p, >_p$ remain unchanged when consumption bundles are scaled. Thus, a data set obeys GAPP if and only if its iso-expenditure version obeys GAPP, which in this case is equivalent to GARP.¹⁵

14. Here, the assumption is that the data span a period over which the mental budget for the observed and unobserved goods is stable. Varying mental budgets would manifest as GAPP violations (see Example 3).

15. There is a similar ‘‘GARP-version’’ of Proposition 1 and that result has been exploited before in the literature (see e.g. Sakai, 1977). Suppose $\mathcal{D} = \{(p^t, x^t)\}_{t=1}^T$ obeys GARP. Then GARP holds even if each observed price vector p^t is arbitrarily scaled. In particular, \mathcal{D} obeys GARP if and only if $\hat{\mathcal{D}} = \{(\hat{p}^t, x^t)\}_{t \in T}$, where $\hat{p}^t = p^t / (p^t \cdot x^t)$, obeys GARP (equivalently, GAPP) since $\hat{p}^t \cdot x^t = 1$ for all $t \in T$. The latter perspective is useful because it highlights the possibility of applying Afriat’s Theorem on $\hat{\mathcal{D}}$, in the *space of prices* (in other words, with the roles of prices and bundles reversed). This immediately gives us a different, ‘‘dual’’ rationalization of \mathcal{D} in terms of indirect utility, that is, there is a continuous, strictly decreasing, and convex function $\tilde{V}: \mathbb{R}_{++}^L \rightarrow \mathbb{R}$ such that $\hat{p}^t \in \arg \min_{\{p \in \mathbb{R}_{++}^L : p \cdot x^t \geq 1\}} \tilde{V}(p)$.




FIGURE 2
Choices that satisfy GAPP but not GARP

Proposition 1. Let $\mathcal{D} = \{(p^t, x^t)\}_{t=1}^T$ be a data set and let $\check{\mathcal{D}} = \{(p^t, \check{x}^t)\}_{t \in T}$, where $\check{x}^t = x^t / (p^t \cdot x^t)$. Then the revealed preference relations \succeq_p^* and \succ_p^* on $\mathcal{P} = \{p^t\}_{t=1}^T$ and the revealed preference relations \succeq_x^* and \succ_x^* on $\check{\mathcal{X}} = \{\check{x}^t\}_{t=1}^T$ are related in the following manner:

- (1) $p^t \succeq_p^* p^{t'}$ if and only if $\check{x}^t \succeq_x^* \check{x}^{t'}$.
- (2) $p^t \succ_p^* p^{t'}$ if and only if $\check{x}^t \succ_x^* \check{x}^{t'}$.

As a consequence, \mathcal{D} obeys GAPP if and only if its iso-expenditure version, $\check{\mathcal{D}}$, obeys GARP.

Proof. Notice that

$$p^t \cdot \frac{x^t}{p^t \cdot x^t} \geq (>) p^{t'} \cdot \frac{x^{t'}}{p^{t'} \cdot x^{t'}} \iff p^{t'} \cdot x^{t'} \geq (>) p^t \cdot x^t.$$

The left side of the equivalence says that $\check{x}^t \succeq_x \check{x}^{t'}$ while the right side says that $p^{t'} \succeq_p p^t$. This implies (1) since \succeq_p^* and \succeq_x^* are the transitive closures of \succeq_p and \succeq_x , respectively.

Similarly, it follows (from the strict inequality version of the above equivalence) that $\check{x}^t \succ_x \check{x}^{t'}$ if and only if $p^t \succ_p p^{t'}$, which leads to (2). The claims (1) and (2) together guarantee that there is a sequence of observations in \mathcal{D} that lead to a GAPP violation if and only if the analogous sequence in $\check{\mathcal{D}}$ lead to a GARP violation. ■

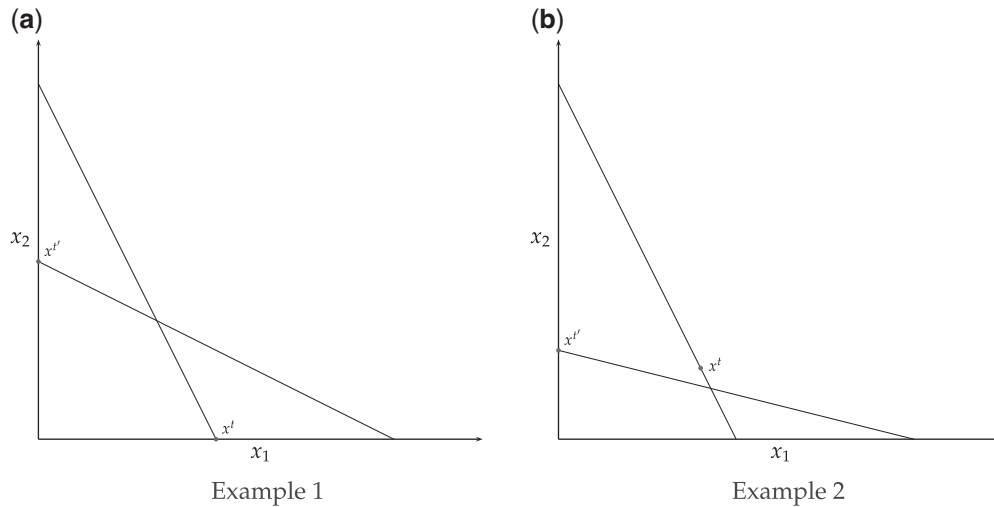


FIGURE 3
Expenditure-normalized choices

As an illustration, compare the data sets in Figures 1 and 2 to the iso-expenditure data sets in Figure 3a and b. It can be clearly observed that the iso-expenditure data in Figure 3a contains a GARP violation (which implies it does not satisfy GAPP) whereas the data in Figure 3b does not violate GARP (and, hence, satisfies GAPP).

Proposition 1 implies that the augmented utility model can be tested in two ways: we can either test GAPP directly or test GARP on its iso-expenditure version. If we are simply interested in testing GAPP on a single-agent data set \mathcal{D} , normalization brings no advantage: the test is computationally straightforward in either case. However, as we shall see in Section 5, iso-expenditure scaling plays an important role in the test we develop (on repeated cross-sectional demand data) for the random utility version of our model.

GARP and GAPP are distinct but not mutually exclusive properties. For instance, data collected from a consumer who maximizes a quasilinear augmented utility satisfies both properties.¹⁶ However, even when both properties are satisfied, demand predictions at an out-of-sample price will differ based on the property employed. Further discussion on the relationship between GAPP and GARP is found in Supplementary Appendix A.1.

Proposition 1 and the fact that scaling consumption bundles does not affect the revealed price preference relation makes it natural to wonder about the relationship between the augmented-utility model and the standard model (as in (3)) with homothetic preferences. A data set that is rationalized in the latter sense (Varian (1983) provides a characterization) has the feature that it must satisfy GARP for *any* arbitrary scaling of consumption bundles and thus will satisfy GAPP. By contrast, a data set that satisfies GAPP must only satisfy GARP for the particular scaling that equalizes expenditure across observations. Thus, GAPP is a less stringent property; that it is *strictly* less stringent is clear from Example 2, which satisfies GAPP but violates GARP and

16. When U has the form (6), x^t maximizes $U(x, -p^t \cdot x)$ only if x^t maximizes $\tilde{U}(x)$ in $\{x \in \mathbb{R}_+^L : p^t \cdot x \leq p^t \cdot x^t\}$. Thus \mathcal{D} must also obey GARP. A broader class of augmented utility functions that satisfy both GAPP and GARP is given in Section A.1.2 of the Supplementary Appendix.

therefore cannot be rationalized in Afriat's sense for any locally non-satiated preference, let alone a homothetic preference.¹⁷

3.2. Preference over prices

We know from Theorem 1 that if \mathcal{D} obeys GAPP then it can be rationalized by an augmented utility function with an indirect utility V that is defined at all price vectors in \mathbb{R}_{++}^L . It is easy to check that any V as defined by (1) has the following properties:

- (a) V is *non-increasing*; that is, if $p' \geq p$ (element by element) then $V(p') \leq V(p)$, and
- (b) V is *quasiconvex*; that is, $V(\beta p + (1 - \beta)p') \leq \max\{V(p), V(p')\}$ for all $\beta \in [0, 1]$.

Any rationalizable data set \mathcal{D} is potentially rationalizable by many augmented utility functions, with each one leading to a different indirect utility. We denote this set of indirect utilities by $\mathbf{V}(\mathcal{D})$. We previously noted that if $p^t \succeq_p^* (>_p^*) p^{t'}$ then $V(p^t) \geq (>) V(p^{t'})$ for any $V \in \mathbf{V}(\mathcal{D})$; in other words, the conclusion that the consumer prefers the prices p^t to $p^{t'}$ is *nonparametric* in that it is independent of the precise augmented utility used to rationalize \mathcal{D} . The next result (proved in Supplementary Appendix A.2) says that, absent further information on the augmented utility, this is *all* the information on the consumer's preference over prices in \mathcal{P} that we can glean from the data. Thus, the revealed price preference relation contains the most detailed information for welfare comparisons.

Proposition 2. *Suppose $\mathcal{D} = \{(p^t, x^t)\}_{t=1}^T$ is rationalizable by an augmented utility function. Then for any $p^t, p^{t'}$ in \mathcal{P} :*

- (1) $p^t \succeq_p^* p^{t'}$ if and only if $V(p^t) \geq V(p^{t'})$ for all $V \in \mathbf{V}(\mathcal{D})$.
- (2) $p^t >_p^* p^{t'}$ if and only if $V(p^t) > V(p^{t'})$ for all $V \in \mathbf{V}(\mathcal{D})$.

3.3. Compensation for a price change

In the standard consumer model, the compensating and equivalent variations are used to quantify the welfare impact of a price change. Analogous concepts exist for the augmented utility model and bounds for them can be recovered from the data.¹⁸

Let U be the consumer's augmented utility function. Suppose, the price changes from p^{t_1} to p^{t_2} leading to a change of x^{t_1} to x^{t_2} in consumption. Then, there exists μ_c such that

$$\max_{x \in \mathbb{R}_+^L} U(x, -p^{t_2} \cdot x - \mu_c) = V(p^{t_1}). \quad (8)$$

Note that μ_c is unique since U is strictly increasing in the last argument. We can think of μ_c as the lump sum transferred from the consumer after the price change that makes her just indifferent between the situation before and after the change.

17. Example A.1 in the Supplementary Appendix contrasts demand predictions using the augmented utility model and the constrained-optimization model (both with and without imposing homotheticity on the preference).

18. Calculations of the compensating and equivalent variations in either the standard model or in ours assume that the preference recovered from consumption data is indeed the consumer's true preference and that is the position we take in this subsection. If behavioural considerations of the type discussed in Section 2.4 are at play, then one could conceptually distinguish between an agent's true preference and that which is recovered from the data under some type of bounded rationality (such as price inattention). The latter preference will still be useful for a positive analysis of demand (e.g. in estimating behaviour at out-of-sample prices) but compensation which is calculated based on such a preference could be problematic.

Suppose we interpret U as arising from an overall utility function $\tilde{U}(x, z)$ (that depends on the observed goods x and the level z of an outside good), given the consumer's wealth of M , so that $U(x, -e) = \tilde{U}(x, M - e)$. Since μ_c solves (8), it will also satisfy

$$\max_{\{x \in \mathbb{R}_+^L : p^{t2} \cdot x \leq M - \mu_c\}} \tilde{U}(x, (M - \mu_c) - p^{t2} \cdot x) = \tilde{U}(x^{t1}, M - p^{t1} \cdot x^{t1}).$$

In other words, μ_c is the reduction in total wealth that will leave the consumer's overall utility at p^{t2} the same as it was at p^{t1} . Thus, with this interpretation of the augmented utility function, μ_c coincides with what is called the *compensating variation* in standard consumer theory and we shall use this term to also refer to μ_c (defined by (8)).

Pushing the analogy further, we can use the compensating variation in our model in the same way it is typically used. For example, a price change from p^{t1} to p^{t2} may benefit some but hurt others. The Kaldor criterion deems this change an overall improvement if the sum of the compensating variations across consumers is positive as it guarantees that those who benefit could, in principle, compensate the losers and still be better off.

In a similar way, we can define the *equivalent variation* as the value μ_e that solves

$$\max_{x \in \mathbb{R}_+^L} U(x, -p^{t1} \cdot x + \mu_e) = V(p^{t2}). \quad (9)$$

If $U(x, -e) = \tilde{U}(x, M - e)$, μ_e coincides with the usual equivalent variation as it solves

$$\max_{\{x \in \mathbb{R}_+^L : p^{t2} \cdot x \leq M + \mu_e\}} \tilde{U}(x, (M + \mu_e) - p^{t1} \cdot x) = \tilde{U}(x^{t2}, M - p^{t2} \cdot x^{t2}).$$

If \mathcal{D} satisfies GAPP and contains observation (p^{t1}, x^{t1}) , what can we say about the compensating variation of a price change from p^{t1} to p^{t2} (where the latter may not be a price observed in \mathcal{D})? Even though there is typically a range of these values since there is more than one augmented utility that rationalizes \mathcal{D} , it is possible to obtain a *tight* lower bound for the set of possible compensating variation values. This is given by

$$\inf\{\mu_c : \mu_c \text{ solves (8) for some augmented utility function } U \text{ that rationalizes } \mathcal{D}\}.$$

Abusing terminology somewhat, we shall denote this term simply by $\inf(\mu_c)$.

We now describe how to compute this bound; we omit the analogous exercise for the equivalent variation. Let $S \subset T$ be the set of observations such that $s \in S$ if $p^s \geq_p^* p^{t1}$. This set is non-empty since it contains p^{t1} itself. For each $s \in S$, there is m_c^s such that

$$p^{t2} \cdot x^s + m_c^s = p^s \cdot x^s. \quad (10)$$

We claim that for any U that rationalizes \mathcal{D} , the compensating variation $\mu_c \geq m_c^s$. This is because if $m < m_c^s$, then $m \neq \mu_c$ for any utility function rationalizing \mathcal{D} . Indeed,

$$\begin{aligned} \max_{x \in \mathbb{R}_+^L} U(x, -p^{t2} \cdot x - m) &\geq U(x^s, -p^{t2} \cdot x^s - m) > U(x^s, -p^{t2} \cdot x^s - m_c^s) \\ &= U(x^s, -p^s \cdot x^s) \geq U(x^{t1}, -p^{t1} \cdot x^{t1}) = V(p^{t1}). \end{aligned}$$

Thus, $\inf(\mu_c) \geq m_c^s$ for all $s \in S$. In fact, it is possible to obtain a stronger conclusion:

$$\inf(\mu_c) = \max\{m_c^s : m_c^s \text{ satisfies (10) for some } s \in S\}. \quad (11)$$

Note that the right side of this equation can be easily computed from the data.

Notice that if p^{t_2} is revealed preferred to p^{t_1} (equivalently, that there is $s' \in S$ such that $m_c^{s'} \geq 0$),¹⁹ then $\inf(\mu_c) \geq 0$; in other words, at $p = p^{t_2}$, a lump sum tax of $\inf(\mu_c)$ will leave the agent no worse off than at t_1 and potentially better off. On the other hand, if p^{t_2} is *not* revealed preferred to p^{t_1} , that is, for every $s \in S$, we have $m_c^s < 0$, then $\inf(\mu_c) < 0$; in other words, at $p = p^{t_2}$, a lump sum transfer of $\inf(\mu_c)$ to the agent will leave the agent no worse off than at t_1 and potentially better off.

Supplementary Appendix A.5 has a fuller discussion and includes a proof of (11).

3.4. Measuring departures from rationality

GARP is frequently violated in empirical applications. The *extent* of such violations for a data set \mathcal{D} is typically measured by the *critical cost efficiency index* Afriat (1973). This is the largest $e \in (0, 1]$ for which there is a locally non-satiated utility function \tilde{U} such that $\tilde{U}(x^t) \geq \tilde{U}(x)$ for all x in the “shrunk” budget set $B_e^t = \{x \in \mathbb{R}_+^L : p^t \cdot x \leq e p^t \cdot x^t\}$. Calculating this index is straightforward via a modified version of GARP (see Afriat, 1973). Rationality is imperfect if $e < 1$ since the consumer behaves as though she ignores bundles x' that satisfy $e p^t \cdot x^t < p^t \cdot x' \leq p^t \cdot x^t$ and, there could be some observation \hat{t} and bundle x' in this range for which $\tilde{U}(x') > \tilde{U}(x^{\hat{t}})$.

We can similarly measure the extent to which a data set \mathcal{D} fails to be rationalized by an augmented utility function. We say that \mathcal{D} is ϑ -rationalized by an augmented utility function if there is an augmented utility $U : \mathbb{R}_+^L \times \mathbb{R}_- \rightarrow \mathbb{R}$ such that, at each t ,

$$U(x^t, -p^t \cdot x^t) \geq U(x, -\vartheta^{-1} p^t \cdot x) \text{ for all } x \in \mathbb{R}_+^L.$$

Note that if \mathcal{D} can be ϑ -rationalized then it can be ϑ' -rationalized for any $\vartheta' < \vartheta$, since U is strictly decreasing in expenditure. A consumer who is ϑ -rational (for $\vartheta < 1$) has limited rationality in the sense that there could be a bundle x' and an observation \hat{t} such that

$$U(x^{\hat{t}}, -p^{\hat{t}} \cdot x') > U(x^{\hat{t}}, -p^{\hat{t}} \cdot x^{\hat{t}}) \geq U(x', -\vartheta^{-1} p^{\hat{t}} \cdot x').$$

In words, the consumer fails to recognize that x' is superior to $x^{\hat{t}}$ at $t = \hat{t}$ because she has inflated (by ϑ^{-1}) the cost of x' . Any data set can be ϑ -rationalized for some $\vartheta \in (0, 1]$ and the supremum ϑ^* over these values provides a natural measure of rationality which we shall refer to as the *rationality index*. The next result (proved in the Supplementary Appendix A.4.3) establishes a connection between this rationality index and Afriat’s efficiency index.

Proposition 3. *Let $\mathcal{D} = \{(p^t, x^t)\}_{t=1}^T$ be a data set and let $\check{\mathcal{D}} = \{(p^t, \check{x}^t)\}_{t \in T}$, where $\check{x}^t = x^t / (p^t \cdot x^t)$, be its iso-expenditure version. Then, ϑ^* is the rationality index for \mathcal{D} if and only if it is the critical cost efficiency index for $\check{\mathcal{D}}$.*

A consequence of this proposition is that the rationality index (like the efficiency index) is easy to compute. Supplementary Appendix A.4.2 discusses its computation in more general settings.

3.5. Deflating prices

For a data set $\mathcal{D} = \{(p^t, x^t)\}_{t=1}^T$ spanning a long period, nominal expenditure may not be an accurate opportunity cost measure due to price changes of both the observed and unobserved goods. This can be accounted for by deflating the prices of the L goods with a general price index. In other

19. Recall that $p^{t_2} \succeq_p^* p^{t_1}$ makes sense even if p^{t_2} is not observed in the data set; see footnote 7.

words, we can check if $\tilde{\mathcal{D}} = \{(p^t/k^t, x^t)\}_{t=1}^T$ obeys GAPP, where $k^t \in \mathbb{R}_{++}$ is an index of the general price level. If it does, it implies that there is an augmented utility function U that rationalizes the data after deflation; in other words,

$$x^t \in \operatorname{argmax}_{x \in \mathbb{R}_+^L} U(x, -(p^t \cdot x)/k^t) \text{ for all } t \in T.$$

This simple adjustment can be precisely justified when the augmented utility is the reduced form of a larger constrained optimization problem. Indeed, suppose that the consumer is maximizing an overall utility $\tilde{U}(x, y)$ subject to $p \cdot x + q \cdot y \leq M$, where x is the observed bundle and y is a bundle of other goods with prices q . Keeping q and M fixed, let $U(x, -e)$ be the greatest overall utility the consumer can achieve by choosing y optimally, subject to expenditure $M - e$ and conditional on consuming x , that is,

$$U(x, -e) = \max\{\tilde{U}(x, y) : y \geq 0 \text{ and } q \cdot y \leq M - e\}. \quad (12)$$

At the prices p^t for the observed goods and q for the outside goods, the consumer chooses a bundle (x^t, y^t) to maximize \tilde{U} subject to $p^t \cdot x + q \cdot y \leq M$. Then, $\mathcal{D} = \{(p^t, x^t)\}_{t=1}^T$ will obey GAPP, since x^t maximizes $U(x, -p^t \cdot x)$, with U as defined by (12).

Now, suppose that the prices of the other goods and global budget are changing and consider the simplest case where they move proportionately: at t , they are $k^t q$ and $k^t M$ for some scalar $k^t > 0$. In other words, the consumer's nominal wealth is keeping pace with price inflation. Then, at t , the consumer maximizes $\tilde{U}(x, y)$ subject to (x, y) obeying $p^t \cdot x + k^t q \cdot y \leq k^t M$. Dividing this inequality by k^t , we see that the consumer's choice is identical to the case where the price of the observed goods is p^t/k^t , with constant external prices q and total wealth M , respectively. Therefore, $\tilde{\mathcal{D}}$ obeys GAPP.

In Supplementary Appendix A.3, we derive a price index (so GAPP holds after deflating p^t) when the relative prices of the outside goods change but this requires stronger assumptions on \tilde{U} .

4. GENERAL CONSUMPTION SPACES AND NON-LINEAR PRICING

We now assume that the consumption space is a set $X \subseteq \mathbb{R}_+^L$, and we define a *price system* as a map $\psi : X \rightarrow \mathbb{R}_+$, where $\psi(x)$ is the cost of purchasing $x \in X$. A special case of a price system is $\psi(x) = p \cdot x$ but the more general formulation with ψ allows for quantity discounts, bundle pricing and other pricing features that can be important in certain contexts (such as our empirical application in Section 7.1). A data set $\mathcal{D} = \{(\psi^t, x^t)\}_{t=1}^T$ consists of price systems and the observed consumption bundles; \mathcal{D} is rationalized by an augmented utility function $U : X \times \mathbb{R}_- \rightarrow \mathbb{R}$ if

$$x^t \in \operatorname{argmax}_{x \in X} U(x, -\psi^t(x)) \text{ for all } t \in T. \quad (13)$$

The notion of revealed preference over prices can be extended to a revealed preference over price systems. We say that $\psi^{t'}$ is *directly revealed preferred* (*directly revealed strictly preferred*) to ψ^t if $\psi^{t'}(x^t) \leq (<) \psi^t(x^t)$; we denote this by $\psi^{t'} \succeq_p (>_p) \psi^t$. We denote the transitive closure of \succeq_p by \succeq_p^* , that is, $\psi^{t'} \succeq_p^* \psi^t$ if there are t_1, t_2, \dots, t_N in T such that $\psi^{t'} \succeq_p \psi^{t_1}$, $\psi^{t_1} \succeq_p \psi^{t_2}$, \dots , $\psi^{t_{N-1}} \succeq_p \psi^{t_N}$, and $\psi^{t_N} \succeq_p \psi^t$; in this case we say that $\psi^{t'}$ is *revealed preferred* to ψ^t . If anywhere along this sequence, it is possible to replace \succeq_p with $>_p$ then we denote that relation by $\psi^{t'} \succ_p^* \psi^t$ and say that $\psi^{t'}$ is *strictly revealed preferred* to ψ^t . It is straightforward to check that if \mathcal{D} can be rationalized by an augmented utility function, then it obeys the following generalization of GAPP to price systems:

there do not exist observations $t, t' \in T$ such that $\psi^{t'} \succeq_p^ \psi^t$ and $\psi^t \succ_p^* \psi^{t'}$.*

The next result shows that the converse is also true. The proof (in fact, of a more general result that allows for errors) is in Supplementary Appendix A.4.1.

Theorem 2. *A data set $\mathcal{D} = \{(\psi^t, x^t)\}_{t=1}^T$ can be rationalized by an augmented utility function if and only if satisfies GAPP.*

Furthermore, suppose that \mathcal{D} satisfies GAPP, X is closed and that, for all $t \in T$, the price systems have the following properties: (i) ψ^t is a continuous function; (ii) for any number M , $\{x \in X : \psi^t(x) \leq M\}$ is a compact set; and (iii) ψ^t is strictly increasing in x_K for some $K \subseteq L$.²⁰ Then, for any closed set $Y \subseteq \mathbb{R}_+^L$ containing X , there is a continuous augmented utility function $U : Y \times \mathbb{R}_- \rightarrow \mathbb{R}$ that rationalizes \mathcal{D} , with $U(x, -e)$ strictly increasing in x_K .

REMARKS: (1) Note that condition (ii) is a weak assumption requiring that there be no arbitrarily large bundles with a bounded price. (2) By definition, an augmented utility function is strictly decreasing in expenditure, but in certain cases it may be natural to require U to be strictly increasing in x_K for some set K (which can be empty). The theorem says that this is possible, so long as the price systems are also strictly increasing in x_K . (3) Note the theorem establishes the additional properties for domains larger than X . We show in Section 4.1 that this is natural in certain applications.

The literature on mental accounting has emphasized the possibility of actors in the economy manipulating the mental budgets of agents. The following example shows how a non-linear GAPP test can be used to detect such phenomena.

Example 3. *A store initially prices two goods at $p^t = (1, 2)$ and a shopper purchases $x^t = (10, 20)$ from the store. The store introduces a scheme where regular customers receive a voucher of 12 dollars for in-store purchases; prices are changed to $p^{t'} = (2, 3/2)$ and the shopper buys $x^{t'} = (20, 20)$.²¹ What is the impact of the voucher?*

Since the value of the voucher is small in terms of total income, the shopper could spread this reward widely across all purchases (including purchases from other stores) and this should result in no (or at least a very small) impact on demand for the store's products. On the other hand, she may have a mental budget for purchases at that store, and the voucher represents an appreciable increase in that mental budget by 12 dollars.

A revealed preference analysis supports the latter hypothesis. If we ignore the voucher, the data are not compatible with the maximization of an augmented utility function since $p^t \cdot x^t = 50 = p^{t'} \cdot x^t$ and $p^{t'} \cdot x^{t'} = 70 > 60 = p^t \cdot x^{t'}$, which violates GAPP. On the other hand, at observation t' , we could model the shopper as mentally discounting 12 dollars from her expenditure at the shop. In formal terms, the price system at t' is a function $\psi^{t'}(x) = \max\{p^{t'} \cdot x - 12, 0\}$, so $\psi^{t'}(x^{t'}) = 58$. In this case, we have $\psi^{t'} \succ_p^ \psi^t$ (where $\psi^t(x) = p^t \cdot x$), but now $\psi^t \not\prec_p^* \psi^{t'}$ since $\psi^t(x^{t'}) = 60 > \psi^{t'}(x^{t'}) = 58$. So the data satisfy GAPP, but with a non-linear pricing system based on the shopper's mental accounting.²²*

4.1. Discrete consumption spaces

Below are three instances where Theorem 2 could be applied.

20. This means that if $x', x \in X$, $x' \neq x$, $x_\ell = x'_\ell$ for all $\ell \notin K$, and $x'_\ell \geq x_\ell$ for all $\ell \in K$, then $\psi^t(x') > \psi^t(x)$.

21. If Good 1 is cheap to procure, this scheme is advantageous to the store, since in the first instance, the shopper spends 50 dollars while in the second, she spends 58 (net of the voucher).

22. Notice (in connection with our discussion of mental accounting in Section 2.4) that the total mental budget of the shopper remains unknown, though the researcher observes an event that has altered that budget.

(1) Suppose that the consumption space $X = \mathbb{N}^K \times \mathbb{R}_+^{L-K}$ (where \mathbb{N} is the set of natural numbers) consists of L goods of which the first K can only be consumed in discrete quantities. Theorem 2 is applicable whether or not prices are linear. In the latter case, the price system is $\psi^t(x) = p^t \cdot x$, which is strictly increasing in x . Theorem 2 guarantees that, if GAPP holds, then there is a continuous augmented utility function that is strictly increasing in x and rationalizes \mathcal{D} .

(2) Another natural environment is one where the consumer purchases a subset of objects from a set with L items. Then each subset can be represented as an element of $X = \{0, 1\}^L$. For $x \in X$, the ℓ th entry x_ℓ equals 1 if and only if the ℓ th object is chosen. If only certain subsets are permissible, then X would be a strict subset of $\{0, 1\}^L$. The price system ψ specifies the cost of different bundles in X . Let e_ℓ denote the vector with 1 in the ℓ th entry and zero everywhere else. Then, $\psi(e_\ell)$ is the price of purchasing good ℓ alone. The price system is non-linear if $\psi(x) \neq \sum_{\ell=1}^L x_\ell \psi(e_\ell)$ for some $x \in X$.

(3) Empirical demand models of differentiated goods typically model goods in terms of their characteristics (see Nevo, 2000). Suppose that there are L characteristics and I goods. Let $Y_\ell \subseteq \mathbb{R}_+$ be the set of values that characteristic ℓ can take. Then, the characteristics space is $Y = \times_{\ell=1}^L Y_\ell$.²³ Each good i has characteristics $x^i \in Y$. Assuming (as is common) that a consumer purchases only one good, the consumption space is $X = \{x^i\}_{i=1}^I$ and a price system $\psi : X \rightarrow \mathbb{R}_{++}$ is just a list of prices for the different goods.

Here, it is natural to model the consumer with an augmented utility function defined on characteristics and expenditures $Y \times \mathbb{R}_+$, even though the products available are only those in X . Furthermore, among the characteristics, there could be those where higher values are unambiguously better, in which case it is natural to require that utility is strictly increasing in those characteristics. Theorem 2 allows for these considerations. If \mathcal{D} obeys GAPP then it can be rationalized by a continuous augmented utility function $U : Y \times \mathbb{R}_+ \rightarrow \mathbb{R}$. Additionally, for a set of characteristics $K \subseteq L$, one could guarantee that $U(y, -e)$ is strictly increasing in y_K so long as $\psi^t(x)$ is strictly increasing in x_K , for all t .

In models of differentiated goods, it is also common to allow for the introduction of new goods and for changes to a product's characteristics.²⁴ Changes to a product's characteristics could potentially lead to a change in the product's utility which, unless taken into account by the test, could lead to a spurious rejection of augmented utility-maximization. When changes to product characteristics are observable, they can be formally captured by allowing the set of alternatives to depend on t ; in Supplementary Appendix A.4.4, we explain how it is possible to modify the GAPP test in Theorem 2 to account for changes of this type.

4.2. Characteristics models with continuous consumption spaces

Now, consider the characteristics space $Y = \mathbb{R}_+^L$, with each product i represented by a vector of characteristics $x^i \in Y$. Goods can be bought in bundles, so the consumption space is the convex cone X generated by $\{x^i\}_{i=1}^I$.²⁵ We assume that the vectors $\{x^i\}_{i=1}^I$ are linearly independent; this guarantees that for each $\hat{x} \in X$, there is a *unique* bundle of goods, $\hat{\alpha} = (\hat{\alpha}_i)_{i=1}^I \in \mathbb{R}_+^I$ such that $\sum_{i=1}^I \hat{\alpha}_i x^i = \hat{x}$. We denote $\hat{\alpha}$ by $\alpha(\hat{x})$. Let $p^t \in \mathbb{R}_{++}^I$ be the prices of the I goods at observation t and so a bundle $x \in X$ costs $\psi^t(x) = p^t \cdot \alpha(x)$.

23. If Characteristic 1 naturally takes on continuous values (such as calories) then we let $Y_\ell = \mathbb{R}_+$. Characteristic 2 could be the brand. Suppose there are two brands, then $Y_2 = \{1, 2\}$, and so on.

24. These changes could be substantive (e.g. a change to a breakfast cereal formula) or it could be a change in advertising expenditure that serves as a proxy for a change in a product's public profile.

25. For a GARP-based test of a model of this type, see Blow, Browning and Crawford (2008).

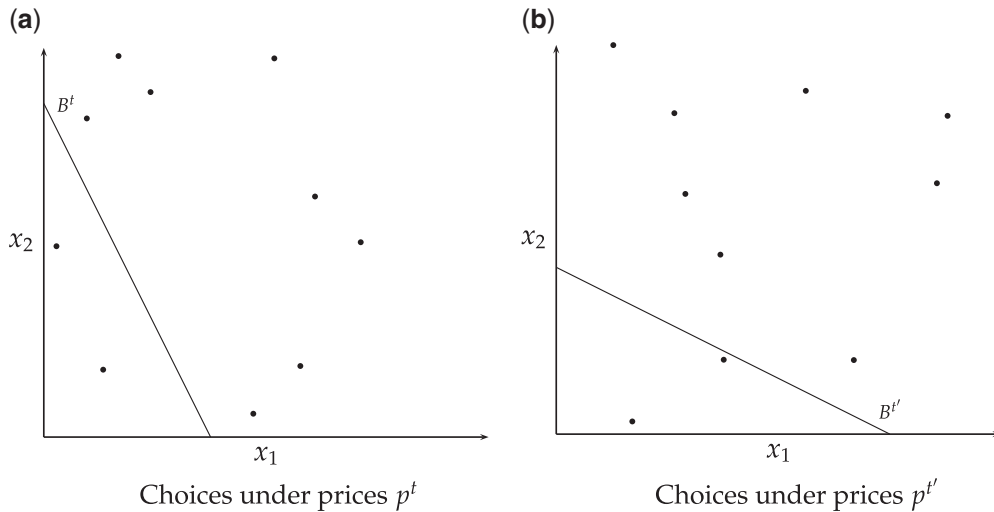


FIGURE 4
The data set

The researcher observes prices p^t and bundles $\alpha^t \in \mathbb{R}_+^I$. We assume $\{x^i\}_{i=1}^I$ is known so the consumption in characteristics space, $x^t = \sum_{i=1}^I \alpha_i^t x^i$, and the price system ψ^t can be imputed. Theorem 2 guarantees that if $\mathcal{D} = \{(\psi^t, x^t)\}_{t \in T}$ satisfies GAPP then it can be rationalized by a continuous augmented utility $U: \mathbb{R}_+^L \times \mathbb{R}_- \rightarrow \mathbb{R}$. If $\psi^t(x)$ is strictly increasing in $x \in X$ for each t , we can also ensure that $U(y, -e)$ is increasing in y .

5. THE RAUM

In this section, we develop the random version of the expenditure-augmented utility model, beginning with an example in which computations can be done in closed form.

5.1. An illustrative example

Example 4. Suppose we have repeated cross-sectional data consisting of the demand of a population of ten consumers at two price vectors. This is illustrated in Figure 4, where the collection of points in the left and right panels indicate the demand bundles at $p^t = (2, 1)$ and $p^{t'} = (1, 2)$, respectively. The lines in Figure 4 merely indicate relative prices.

As this is a repeated cross-section, we cannot match consumption bundles across the two panels by consumer identity. The question is whether this data set can be rationalized, by which we mean the following:

Can the choices at t and t' be matched to form ten distinct pairs such that each pair is rationalized by an augmented utility function (or, equivalently, satisfies GAPP)?

The interpretation is that each pair of choices corresponds to the demand of a single consumer and so rationalization requires the existence of a (time-invariant) distribution of consumer types or, equivalently, a random augmented utility. For more conventional models of utility maximization, this question was analysed by McFadden and Richter (1991, for discrete choice) and McFadden




FIGURE 5
Observed and rescaled choices

(2005, for linear budgets). While these settings appear quite different from ours, the tight link established in Proposition 3 allows us to build on these results.

Specifically, a pair $\mathcal{D} = \{(p^t, x), (p^{t'}, x')\}$ created by choosing bundle x from observation t and x' from observation t' obeys GAPP if and only if its iso-expenditure analogue, $\tilde{\mathcal{D}} = \{(p^t, \tilde{x}), (p^{t'}, \tilde{x}')\}$ as defined in Proposition 3, obeys GARP. This is visualized in Figures 5a and b: replacing data points by iso-expenditure analogues is equivalent to projecting them along origin rays onto the budget corresponding to unit expenditure. Figure 5c superimposes the scaled bundles from both cross-sections. Rationalizability by a random augmented utility is therefore equivalent to asking whether these rescaled observations can be rationalized by a standard RUM (which in this simple case means sorting them into pairs of observations that each obey GARP).

In this example, $\tilde{\mathcal{D}}$ satisfies GARP if, and only if, it is not the case that $\tilde{x} \in B^{2,t}$ and $\tilde{x}' \in B^{1,t'}$, where $(B^{2,t}, B^{1,t'})$ are indicated in Figure 5c. Instead $\tilde{\mathcal{D}}$ must be one of three individually

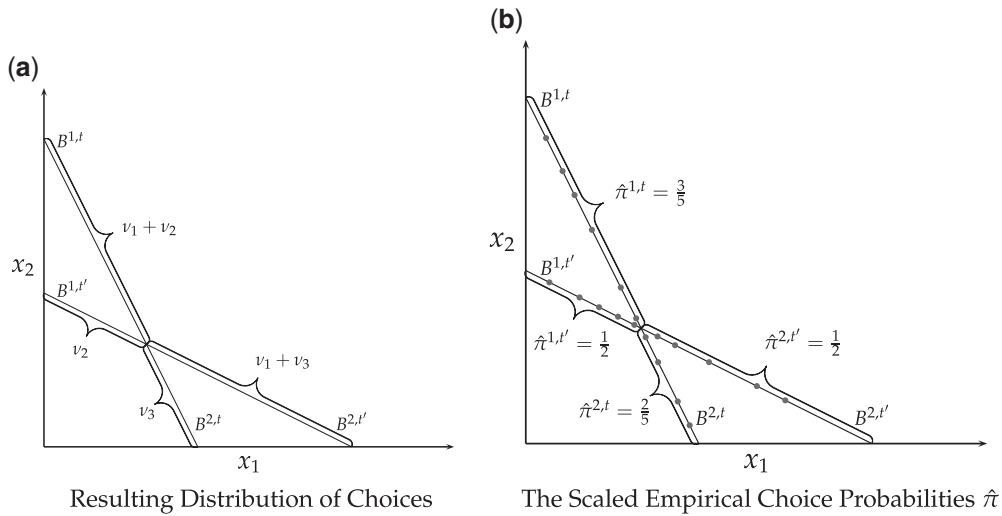


FIGURE 6
Choice distribution and empirical frequency

rationalizable choice types: either $(\check{x}, \check{x}') \in B^{1,t} \times B^{2,t'}$ (no revealed preference), or $(\check{x}, \check{x}') \in B^{1,t} \times B^{1,t'}$ (\check{x} revealed preferred to \check{x}'), or $(\check{x}, \check{x}') \in B^{2,t} \times B^{2,t'}$ (\check{x}' revealed preferred to \check{x}).

Denote the population share of these types by (v_1, v_2, v_3) . They must generate the observed proportion of choices on the segments $(B^{1,t}, B^{2,t}, B^{1,t'}, B^{2,t'})$. Figure 6a visualizes how observed choice probabilities relate to (v_1, v_2, v_3) ; Figure 6b gives the corresponding sample proportions

$$\hat{\pi} = (\hat{\pi}^{1,t}, \hat{\pi}^{1,t'}, \hat{\pi}^{2,t}, \hat{\pi}^{2,t'})' = \left(\frac{3}{5} \frac{2}{5} \frac{1}{2} \frac{1}{2} \right)'. \tag{14}$$

The empirical choice frequencies are rationalizable by a random augmented utility if

$$v_1 + v_2 = \hat{\pi}^{1,t}, \quad v_1 + v_3 = \hat{\pi}^{2,t'}, \quad v_2 = \hat{\pi}^{1,t'}, \quad v_3 = \hat{\pi}^{2,t} \tag{15}$$

can be solved for non-negative (v_1, v_2, v_3) .

This is indeed the case in the example, where the equations are uniquely solved by $(v_1, v_2, v_3)' = (\frac{1}{10}, \frac{1}{2}, \frac{2}{5})'$.²⁶ To further confirm this, we could pair up the bundles on the two budget lines in order of the consumption of x_2 . Then, it is easily seen that each pair satisfies GARP (and hence the corresponding un-scaled pairs satisfy GAPP).

5.2. Rationalization by random augmented utility

Consider now a repeated cross-sectional data set, $\mathcal{D} := \{(p^t, \hat{\pi}^t)\}_{t=1}^T$, where each observation consists of a price p^t and a probability measure $\hat{\pi}^t$ on \mathbb{R}_+^L representing the demand distribution in the population at that price. We now provide a general definition of rationalizability for such a data set.

26. In general, the solution is not unique. In this simple example, it is straightforward to check that the data can be rationalized by a random augmented utility if, and only if, $\hat{\pi}^{1,t} + \hat{\pi}^{2,t'} \geq 1$.

Definition 5.1. *The repeated cross-sectional data set $\mathcal{D} = \{(p^t, \hat{\pi}^t)\}_{t=1}^T$ is rationalized by the RAUM if there exists a probability space $(\Omega, \mathcal{F}, \mu)$ and a random variable $\chi : \Omega \rightarrow (\mathbb{R}_+^L)^T$ such that, almost surely, $\{(p^t, \chi^t(\omega))\}_{t \in T}$ can be rationalized by an augmented utility function (equivalently, obeys GAPP) and*

$$\hat{\pi}^t(Y) = \mu(\{\omega \in \Omega : \chi^t(\omega) \in Y\}) \text{ for any measurable } Y \subseteq \mathbb{R}_+^L. \quad (16)$$

In this definition, one could interpret Ω as the population of consumers and $\chi^t(\omega)$ as the demand of consumer type ω at price p^t . All consumer types in the support of μ must be consistent with the augmented utility model and, for all t , the observed distribution of demand $\hat{\pi}^t$ must coincide with that induced by the distribution μ over consumer types. Alternatively, individuals' augmented utility functions might change over time but in such a way that the population distribution is stationary.

In Example 4, the data set contains two cross-sectional distributions, both of which are discrete with 10 mass points. A RAUM-rationalization involves matching observations in t with those in t' , so that each pair obeys GAPP. In the general case with T cross-sections, the function χ solves a T -fold matching problem, where each group $\{\chi^t(\omega)\}_{t \in T}$ (along with the associated prices) satisfies GAPP and agrees with the observations (that is, (16) is satisfied).²⁷

Theorem 3 below characterizes the rationalizability of $\mathcal{D} = \{(p^t, \hat{\pi}^t)\}_{t=1}^T$. It is proved in Supplementary Appendix, but we now explain it heuristically. Let $B^t := \{x \in \mathbb{R}_+^L : p^t \cdot x = 1\}$ be the budget plane at prices p^t and expenditure 1 and let $\check{\pi}^t$ be the distribution that obtains after projecting $\hat{\pi}^t$ onto B^t .²⁸ We refer to $\check{\mathcal{D}} = \{(p^t, \check{\pi}^t)\}_{t=1}^T$ as the iso-expenditure analogue of \mathcal{D} . We say that $\check{\mathcal{D}}$ is rationalized by the RUM if there is a probability space $(\Omega, \mathcal{F}, \mu)$ and a random variable $\chi : \Omega \rightarrow (\mathbb{R}_+^L)^T$ such that, almost surely, $\{(p^t, \chi^t(\omega))\}_{t \in T}$ obeys GARP and $\check{\pi}^t(Y) = \mu(\{\omega \in \Omega : \chi^t(\omega) \in Y\})$ for any measurable $Y \subseteq \mathbb{R}_+^L$. Crucially, reasoning along the lines of Proposition 3 establishes that \mathcal{D} can be RAUM-rationalized if, and only if, $\check{\mathcal{D}} = \{(p^t, \check{\pi}^t)\}_{t=1}^T$ can be RUM rationalized. Finally, to check the latter, we simply adopt the procedure laid out in McFadden (2005) and KS for testing RUM on iso-expenditure data sets, which we now explain.

For ease of exposition, we impose the following assumption.

Assumption 1. *For all $t, t' \in T$ with $B^t \neq B^{t'}$,*

$$\hat{\pi}^t \left(\left\{ x \in \mathbb{R}_+^L : \frac{x}{p^t \cdot x} \in B^t \text{ and } \frac{x}{p^{t'} \cdot x} \in B^{t'} \right\} \right) = 0.$$

This assumption excludes (with probability 1) choices on the intersection of budget planes. It is not required for any of our results but simplifies the exposition because it forces revealed preferences to be strict.²⁹ It is always satisfied if $\hat{\pi}^t$ is absolutely continuous with respect to Lebesgue measure and is unlikely to be violated in any application with a continuous consumption space and linear prices.

Next, for any budget B^t , let $\{B^{1:t}, \dots, B^{t:t}\}$ denote the collection of subsets that each subset has as its boundaries the intersection of B^t with other budget sets and/or the boundary

27. It is straightforward to check that, with two observations, finding a rationalization is equivalent to finding a zero-cost solution to the transportation problem (see Galichon and Henry, 2011) where the cost of a pair of bundles is 0 if it obeys GAPP and 1 otherwise.

28. Formally, given a measurable set C in B^t , $\check{\pi}^t(C) = \hat{\pi}^t(\bar{C})$, where \bar{C} is the cone generated by C .

29. If we allow for mass at budget intersections, then we would have to include them in our definition of patches. This is notationally cumbersome but once included our arguments (and Theorem 3) remain correct.

planes of the positive orthant. These are the higher-dimensional and multi-period analogues to the line segments in Figure 5c. Formally, for all $t \in T$ and $i_t \neq i'_t$, each set in $\{B^{1,t}, \dots, B^{T,t}\}$ is closed and convex and satisfies the following conditions:

- (i) $\cup_{1 \leq i_t \leq I_t} B^{i_t,t} = B^t$,
- (ii) $\text{int}(B^{i_t,t}) \cap B^{t'} = \emptyset$ for all $t' \neq t$ that satisfy $B^t \neq B^{t'}$ (where $\text{int}(B^{i_t,t})$ denotes the relative interior of $B^{i_t,t}$),
- (iii) $B^{i_t,t} \cap B^{i'_t,t} \neq \emptyset$ implies that $B^{i_t,t} \cap B^{i'_t,t} \subset B^{t'}$ for some $t' \neq t$ that satisfies $B^t \neq B^{t'}$.

We will henceforth call these sets *patches*. For each patch $B^{i_t,t}$, let

$$\pi^{i_t,t} := \overset{\circ}{\pi}^t \left(\left\{ x \in \mathbb{R}_+^L : \frac{x}{p^t \cdot x} \in B^{i_t,t} \right\} \right) \tag{17}$$

be the probability that a period- t choice after scaling lies on patch $B^{i_t,t}$. Denote by π^t the vector $(\pi^{i_t,t})_{i_t=1}^{I_t}$ and by π the column vector $(\pi^1, \pi^2, \dots, \pi^T)'$ of *observed patch probabilities*. (Assumption 1 here causes simplification because it guarantees that $\sum_{i_t=1}^{I_t} \pi^{i_t,t} = 1$.) These patch probabilities are relevant because, as we shall explain, rationalizability depends only on the patch probabilities and not on the distribution within each patch.

To any deterministic, iso-expenditure data set $\check{D} = \{(p^t, \check{x}^t)\}$, we can associate a vector $a = (a^{1,1}, \dots, a^{T,T})$, where $a^{i_t,t} = 1$ if $\check{x}^t \in B^{i_t,t}$ and 0 otherwise. The first crucial observation is that any two deterministic iso-expenditure data sets that are represented by the same vector a would either *both* obey or fail GARP (because their revealed preference relations have the same structure).³⁰ Thus, even though there are infinitely many deterministic iso-expenditure data sets \check{D} that obey GARP, they belong to a *finite* set of equivalence classes, each of which is represented by a vector a . We gather these vectors into the set \mathcal{A} . The second crucial observation is that $\check{\mathcal{D}} = \{(p^t, \check{\pi}^t)\}_{t=1}^T$ can be RUM-rationalized if, and only if, π is in the convex hull of vectors in \mathcal{A} (so π is generated by a distribution over GARP-consistent types).

To state this a bit more formally, collect all distinct GARP-consistent vectors $a \in \mathcal{A}$ into the columns of a matrix A . In Example 4, this matrix equals

$$A = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}, \tag{18}$$

with column j representing type j ($j = 1, 2, 3$).³¹ Then the cross-sectional data set \mathcal{D} can be RAUM-rationalized if, and only if, its iso-expenditure analogue $\check{\mathcal{D}}$ can be RUM-rationalized and the latter holds if, and only if, there is $v_a \geq 0$ (corresponding to weight of type $a \in \mathcal{A}$), such that $\sum_{a \in \mathcal{A}} v_a = 1$ and

$$\pi^{i_t,t} = \sum_{a \in \mathcal{A}} v_a a^{i_t,t} \tag{19}$$

for all $B^{i_t,t}$. More succinctly, $Av = \pi$, where v is the column vector $(v_a)_{a \in \mathcal{A}}$.³² The next theorem summarizes our discussion.

30. This observation is made in Section 3.1 of McFadden (2005).
 31. Each column has four entries because there are four patches in total. For example, the first column represents type 1 data sets $\check{D} = \{(p^1, \check{x}), (p^1, \check{x}')\}$, where $(\check{x}, \check{x}') \in B^{1,1} \times B^{2,1}$ (see Figure 5c).
 32. In Example 4, A and π are given by (18) and (14), respectively.

Theorem 3. Let $\mathcal{D} = \{(p^t, \hat{\pi}^t)\}_{t=1}^T$ be a repeated cross-sectional data set obeying Assumption 1. Then, \mathcal{D} can be RAUM-rationalized if and only if there exists a $v \in \Delta^{|\mathcal{A}|-1}$ such that $Av = \pi$.

It is worth reiterating that even though the RAUM-test given by this theorem requires the (straightforward) calculation of π from an iso-expenditure analogue of \mathcal{D} , the data set \mathcal{D} need not be an iso-expenditure data set. In contrast, the RUM-test set out in McFadden (2005) requires \mathcal{D} itself to be an iso-expenditure data set. Of course, actual data generated by heterogeneous consumers will typically not be iso-expenditure, and indeed the Family Expenditure Survey data used by KS in their implementation of the RUM-test are not. For this reason, their empirical analysis starts by *estimating* an iso-expenditure data set (corresponding to the median expenditure), before implementing the RUM-test on this estimated data set. Obviously, this requires additional econometric work and, therefore, introduces both more assumptions and (albeit asymptotically negligible) noise.

5.3. Welfare comparisons

Since the test for rationalizability involves finding a distribution v over different types, it is possible to use this distribution for welfare analysis. To be specific, suppose that a government is contemplating a change in sales tax that could lead to prices changing from its current value $p^{t'}$ to \hat{p} . Relevant to the government's re-election prospects is the proportion of consumers who will be better off as a result of this price change.³³ Our methods allow us to obtain information on this proportion.

So consider a data set \mathcal{D} that contains among its observations the prevailing prices $p^{t'}$ and the demand distribution $\hat{\pi}^{t'}$. To determine the welfare effect of a price change from $p^{t'}$ to \hat{p} , let $\mathbb{1}_{\hat{p} \succeq_p^* p^{t'}}$ denote the row vector with its length equal to the number of rational types ($|\mathcal{A}|$), such that the j th element is 1 if $\hat{p} \succeq_p^* p^{t'}$ for the rational type corresponding to column j of A and 0 otherwise.³⁴ In words, $\mathbb{1}_{\hat{p} \succeq_p^* p^{t'}}$ enumerates the set of rational types for which \hat{p} is revealed preferred to $p^{t'}$. If \mathcal{D} is rationalizable, Theorem 3 guarantees that

$$\underline{\mathcal{N}}_{\hat{p} \succeq_p^* p^{t'}} := \min_v \mathbb{1}_{\hat{p} \succeq_p^* p^{t'}} v, \quad \text{subject to } Av = \pi, \quad (20)$$

is the lower bound on the proportion of consumers who are revealed better off at prices \hat{p} compared to $p^{t'}$, while the upper bound is

$$\overline{\mathcal{N}}_{\hat{p} \succeq_p^* p^{t'}} := \max_v \mathbb{1}_{\hat{p} \succeq_p^* p^{t'}} v, \quad \text{subject to } Av = \pi. \quad (21)$$

Since (20) and (21) are both linear programs (which have solutions if, and only if, \mathcal{D} is rationalizable), they are easy to implement in practice. Suppose that the solutions are \underline{v} and \overline{v} respectively; then for any $\beta \in [0, 1]$, $\beta \underline{v} + (1 - \beta) \overline{v}$ is also a solution to $Av = \pi$ and, in this case, the proportion of consumers who are revealed better off at \hat{p} compared to $p^{t'}$ is exactly $\beta \underline{\mathcal{N}}_{\hat{p} \succeq_p^* p^{t'}} + (1 - \beta) \overline{\mathcal{N}}_{\hat{p} \succeq_p^* p^{t'}}$. In other words, the proportion of consumers who are revealed better off can take any value in the interval $[\underline{\mathcal{N}}_{\hat{p} \succeq_p^* p^{t'}}, \overline{\mathcal{N}}_{\hat{p} \succeq_p^* p^{t'}}]$.

Proposition 2 tells us that the revealed preference relations are tight, in the sense that if, for some consumer, \hat{p} is not revealed preferred to $p^{t'}$ then there exists an augmented utility function

33. We would like to thank an anonymous referee for suggesting this motivation.

34. Even though \hat{p} is not among the observed prices, one could still define $\hat{p} \succeq_p^* p^{t'}$; see footnote 7.

which rationalizes her consumption choices and for which she strictly prefers $p^{t'}$ to \hat{p} . Given this, we know that, amongst all rationalizations of \mathcal{D} , $\underline{N}_{\hat{p} \succeq_p^* p^{t'}}$ is also the infimum on the proportion of consumers who are better off at \hat{p} compared to $p^{t'}$.

The following proposition summarizes these observations.

Proposition 4. *Let $\mathcal{D} = \{(p^t, \hat{\pi}^t)\}_{t=1}^T$ be a repeated cross-sectional data set that satisfies Assumption 1 and is rationalized by the RAUM. Then, for every $\eta \in [\underline{N}_{\hat{p} \succeq_p^* p^{t'}}, \overline{N}_{\hat{p} \succeq_p^* p^{t'}}]$, there is a rationalization of \mathcal{D} for which η is the proportion of consumers who are revealed better off at \hat{p} compared to $p^{t'}$. Furthermore, $\underline{N}_{\hat{p} \succeq_p^* p^{t'}}$ is the infimum of the proportion of consumers who are better off at \hat{p} compared to $p^{t'}$, among all the rationalizations of \mathcal{D} .*

It is helpful to apply Proposition 4 to Example 4. There, the solution to $Av = \pi$ is unique. Of the three types discussed in Section 5.1, only the second one reveals $p^t \succeq_p^* p^{t'}$, so the proportion of consumers revealed better off at p^t compared to $p^{t'}$ equals v_2 . Formally, we have $\mathbb{1}_{p^t \succeq_p^* p^{t'}} = (0, 1, 0)$, $\mathbb{1}_{p^t \succeq_p^* p^{t'}} \cdot v = 1/2$, and $\underline{N}_{p^t \succeq_p^* p^{t'}} = \overline{N}_{p^t \succeq_p^* p^{t'}} = 1/2$. By similar reasoning, we have $\underline{N}_{p^{t'} \succeq_p^* p^t} = \overline{N}_{p^{t'} \succeq_p^* p^t} = 2/5$. The point identification of these quantities is due to the uniqueness of v , which is specific to that simple example.

6. STATISTICAL TEST OF RAUM AND INFERENCE FOR COUNTERFACTUALS

This section outlines our econometric methodologies. First, Section 6.1 provides a statistical test of the RAUM (presented in Section 5). Second, and more importantly, Section 6.2 develops a new method for obtaining asymptotically uniformly valid confidence intervals for counterfactual objects. This result applies to a general class of RUMs, including the RAUM. It can be used for statistical analyses of welfare comparisons, and we use it for that purpose in our empirical study in Section 7.2.

6.1. Testing the RAUM

Recall from Theorem 3 that, given a set of prices and corresponding demand distributions $\mathcal{D} = \{(p^t, \hat{\pi}^t)\}_{t=1}^T$ and an implied vector π of choice probabilities on rescaled and discretized budgets, a test of the RAUM is a test of

$$H_0: \exists v \in \Delta^{|\mathcal{A}|-1} \text{ such that } Av = \pi \iff \min_{v \in \mathbb{R}_+^{|\mathcal{A}|}} [\pi - Av]' \Omega [\pi - Av] = 0, \tag{22}$$

where Ω is a positive definite matrix. The equivalence was noted and exploited in KS.³⁵

In practice, we estimate π by its sample analogue $\hat{\pi} = (\hat{\pi}^1, \dots, \hat{\pi}^T)$ obtained by rescaling the empirical distribution of choices $\{x_{n_t}^t\}_{n_t=1}^{N_t}$, where N_t is the number of observed choices in the data in period t . This gives rise to test statistic

$$J_N := N \min_{v \in \mathbb{R}_+^{|\mathcal{A}|}} [\hat{\pi} - Av]' \Omega [\hat{\pi} - Av], \tag{23}$$

35. The strategy to configure H_0 as a quadratic program also appears in De Paula, Richards-Shubik and Tamer (2018), albeit for a different program and in a different context.

where $N = \sum_{t=1}^T N_t$ denotes the total number of observations. We use the modified bootstrap procedure of KS to compute critical values for this test.

6.2. Inference for counterfactuals in a general class of RUMs

A counterfactual quantity in a RUM can be generally regarded as a function of the underlying distribution v of individual preferences. This section focuses on the case where this mapping is linear, so that we are concerned with statistical inference for $\theta = \rho \cdot v$, where $\rho \in \mathbb{R}^{|\mathcal{A}|}$ is a known vector which varies with the counterfactual of interest. Our analysis of welfare comparisons in Section 5.3 falls into this framework, by letting θ be the proportion of consumers who are revealed better off at prices \hat{p} compared to p^i , with $\rho = \mathbb{1}_{\hat{p} \geq_p^* p^i}$. It is worth emphasizing that the methodology developed in this section has broad applicability: it can be used to study other RUMs (such as the model in Kitamura and Stoye (2019)) and to investigate other objects of interest in RUMs; for example, Lazzati *et al.* (2018) applies our technique to estimate the proportion of non-strategic players in a game.

Note that θ is partially identified as follows:

$$\theta \in \Theta_I \quad \text{where} \quad \Theta_I := \{\rho \cdot v \mid v \geq 0, Av = \pi\}.$$

Our confidence interval inverts a test of

$$\pi \in \mathcal{S}(\theta) \quad \text{where} \quad \mathcal{S}(\theta) := \left\{ Av \mid \rho \cdot v = \theta, v \in \Delta^{|\mathcal{A}|-1} \right\} \quad (24)$$

or equivalently,

$$\min_{v \in \Delta^{|\mathcal{A}|-1}, \theta = \rho \cdot v} [\pi - Av]' \Omega [\pi - Av] = 0.$$

The test statistic is a scaled sample analogue

$$J_N(\theta) = N \min_{v \in \Delta^{|\mathcal{A}|-1}, \theta = \rho \cdot v} [\hat{\pi} - Av]' \Omega [\hat{\pi} - Av] = N \min_{\eta \in \mathcal{S}(\theta)} [\hat{\pi} - \eta]' \Omega [\hat{\pi} - \eta], \quad (25)$$

where the second equality follows from (24). The naive bootstrap fails to deliver valid critical values for (25) as its asymptotic distribution changes discontinuously, depending on the location of π relative to the polytope $\mathcal{S}(\theta)$. A simple application of the modified bootstrap algorithm in KS does not work, as their method relies on, among other things, the polytope $\{Av : v \geq 0\}$ being a cone. This is not necessarily the case for counterfactual analysis, and we need to deal with $\mathcal{S}(\theta)$ without relying on conical properties.

That said, as in KS, we do gain an insight from Weyl–Minkowski duality. In Supplementary Appendix A.7, we show that there exist non-stochastic matrices B, \tilde{B} and a non-stochastic vector-valued function $d(\theta)$ such that $\pi \in \mathcal{S}(\theta)$ if, and only if,

$$B\pi \leq 0, \tilde{B}\pi = d(\theta) \quad \text{and} \quad \mathbf{1} \cdot \pi = 1, \quad (26)$$

where $\mathbf{1}$ is the I -vector of ones where $I = \sum_{t=1}^T I_t$ is the total number of patches. Thus, in principle, this is a linear (in)equality testing problem. There is a rich literature on such problems. However, we cannot directly invoke that literature because we cannot compute (B, \tilde{B}) in practice for a problem with a relevant scale.

While we therefore need to work with representation (24), representation (26) is useful. It illustrates that the inference problem is non-standard; in particular, the limiting distribution of the

test statistic depends on how close to binding each of the constraints encoded in $(B, \tilde{B}, d(\theta))$ is. From analogy to the moment inequalities literature, it also pretty much implies that the constraints' slackness cannot be pre-estimated with sufficient accuracy; the reason being that it enters the test's asymptotic representation scaled by \sqrt{N} . However, we also know that certain existing procedures which shrink the estimated slack of all inequalities to zero before computing the distribution of J_N will work. Our proposal is inspired by these but must implement the idea with the computationally feasible representation (24) instead of (26), which is only theoretically available. This means that we cannot calculate the empirical slack, which is explicit in (the empirical version of) representation (26) but not in (24), which is why a new method is called for.

Intuitively, we contract (or "tighten") the polytope $\mathcal{S}(\theta)$ toward a point in its relative interior, thereby effectively (but non-obviously) reducing the empirical slack in any inequality constraint. This forces all the constraints with small slacks to be binding after "tightening". Note that, unlike in KS, we face substantial added complications because (i) we need to deal with a non-conical $\mathcal{S}(\theta)$, and (ii) the appropriate way to tighten the polytope $\mathcal{S}(\theta)$ varies with the value of θ through the dependence of $\mathcal{S}(\theta)$ on θ . This leads to a *restriction-dependent tightening* approach which we now describe in broad strokes.

Choose a sequence τ_N such that $\tau_N \downarrow 0$ and $\sqrt{N}\tau_N \uparrow \infty$ (we make a specific proposal in the appendix) and define

$$\mathcal{S}_{\tau_N}(\theta) := \{A\nu \mid \rho \cdot \nu = \theta, \nu \in \mathcal{V}_{\tau_N}(\theta)\},$$

where $\mathcal{V}_{\tau_N}(\theta)$ is obtained by appropriately constricting $\Delta^{|\mathcal{A}|-1}$; in particular, some components of ν are forced to be boundedly above 0. Note that $\mathcal{S}_{\tau_N}(\theta)$ depends on θ through the equation $\rho \cdot \nu = \theta$ but also because, as the notation suggests, the construction of $\mathcal{V}_{\tau_N}(\theta)$ will change with θ , a key feature of our algorithm. The definition of $\mathcal{V}_{\tau_N}(\theta)$ for general ρ is rather involved and thus deferred to the Supplementary Appendix A.7, but it considerably simplifies for binary ρ as in our application.

The set $\mathcal{S}_{\tau_N}(\theta)$ replaces $\mathcal{S}(\theta)$ in the bootstrap population. The precise algorithm proceeds as follows. For each $\theta \in \Theta$:

- (i) Compute the τ_N -tightened restricted estimator of the empirical choice distribution

$$\hat{\eta}_{\tau_N} := \operatorname{argmin}_{\eta \in \mathcal{S}_{\tau_N}(\theta)} N[\hat{\pi} - \eta]' \Omega[\hat{\pi} - \eta].$$

- (ii) Define the τ_N -tightened recentered bootstrap estimators

$$\hat{\pi}_{\tau_N}^{*(r)} := \hat{\pi}^{*(r)} - \hat{\pi} + \hat{\eta}_{\tau_N}, \quad r = 1, \dots, R,$$

where $\hat{\pi}^{*(r)}$ is a bootstrap analogue of $\hat{\pi}$, and R is the number of bootstrap samples. For instance, in our application, $\hat{\pi}^{*(r)}$ is generated by the simple non-parametric bootstrap of choice frequencies.

- (iii) For each $r = 1, \dots, R$, compute

$$J_{N, \tau_N}^{*(r)}(\theta) = \min_{\eta \in \mathcal{S}_{\tau_N}(\theta)} N[\hat{\pi}_{\tau_N}^{*(r)} - \eta]' \Omega[\hat{\pi}_{\tau_N}^{*(r)} - \eta].$$

- (iv) Use the empirical distribution of $J_{N, \tau_N}^{*(r)}(\theta)$ to obtain the critical value for $J_N(\theta)$.

A confidence interval for θ collects values of θ that are not rejected.

Theorem 4 below (proved in the Supplementary Appendix A.7) establishes asymptotic validity of the above procedure. Let

$$\mathcal{F} := \{(\theta, \pi) \mid \theta \in \Theta, \pi \in \mathcal{S}(\theta) \cup \mathcal{P}\},$$

where \mathcal{P} denote the set of all π that satisfy the technical Condition 1 in the Supplementary Appendix A.7.

Theorem 4. *Choose τ_N so that $\tau_N \downarrow 0$ and $\sqrt{N}\tau_N \uparrow \infty$. Also, let Ω be diagonal. Then under Assumptions 2 and 3 stated in the Supplementary Appendix A.7,*

$$\liminf_{N \rightarrow \infty} \inf_{(\theta, \pi) \in \mathcal{F}} \Pr\{J_N(\theta) \leq \hat{c}_{1-\alpha}\} = 1 - \alpha,$$

where $0 \leq \alpha \leq \frac{1}{2}$ and $\hat{c}_{1-\alpha}$ is the $1 - \alpha$ quantile of J_{N, τ_N}^* .

Note that we have assumed the prices are exogenous throughout our analysis in this section. If the exogeneity condition holds *conditional* on some observable covariates, it is straightforward to incorporate it by replacing the sample analogue $\hat{\pi}$ of π with an appropriate conditional choice probability estimator. Or, in situations where the control function approach is applicable, we can deal with endogeneity along the line of analysis in KS (see their Theorem 5.2). On the other hand, a satisfactory treatment of, for example, price endogeneity caused by unobserved characteristics calls for further extension of our approach. If an appropriate instrumental variable is available, then it might be possible to generalize the “ \mathcal{V} -representations” of identified sets to accommodate it, though we leave such analysis to future research.

7. EMPIRICAL APPLICATIONS

We now present two separate applications meant to show how both the deterministic and random versions of our model can be tested and employed for welfare analysis.

7.1. *Augmented utility model: testing and welfare analysis on Progresa data*

We apply the deterministic model to the Progresa-Oportunidades data set, a workhorse of the treatment evaluation literature. Progresa was a conditional cash transfer program aimed at poor communities in Mexico. The program was remarkable in that it was rolled out in random order so the causal effect of the cash transfers could be studied. For brevity, we do not describe the program in detail; information on the program is widely available including in the paper we discuss next.

Our application builds on recent work of Attanasio and Pastorino (2020) (henceforth AP) who analyse whether the program led to changes in the market prices for basic staples: rice, kidney beans, and sugar. This is an important question because the welfare effect of these transfers would clearly depend in part on their impact on prices. While the previous literature had documented that *average* prices were not affected by the program Hodinott, Skoufias and Washburn (2000), AP argue that sellers charge non-linear prices and that these non-linear price schedules had changed.

Because treatment was randomized across villages but means-tested at the household level, some households faced a changing price schedule but no shock to their own income. In our study, we focus our attention on these households because we can be more confident that their augmented utility functions are unchanged across the observation periods. Our objectives are, firstly, to test the augmented utility model and, secondly, to evaluate the welfare impact of price

changes using that model. This data set is well suited for analysis using our deterministic model because its panel structure means that we can study each household separately. Following AP, we consider non-linear prices, which allows us to implement the results in Section 4.

The theoretical part of AP derives the optimal (non-linear) pricing schedule under the assumption that there is a heterogenous population of households, each of which maximizes a quasilinear utility function subject to a subsistence constraint. This constraint requires a household to consume a minimum number of calories which can be obtained from either the observed bundle x or the numeraire; given x , $\underline{z}(x)$ denotes the minimum amount of the numeraire good needed to meet the calorie threshold. Thus, the household can only choose bundles x that satisfy $\psi(x) + \underline{z}(x) \leq M$, where ψ is the price system and M is household wealth. It is worth noting that the augmented utility framework is sufficiently flexible to accommodate this behaviour. Indeed, the household could be thought of as maximizing an augmented utility function of the modified-quasilinear form

$$U(x, -e) = \tilde{U}(x) - \mathbf{K}(e + \underline{z}(x) - M)e,$$

where $\mathbf{K}(w) = 1$ if $w \leq 0$ and $\mathbf{K}(w)$ is a very large positive number if $w > 0$. In this way, any $(x, -e)$ (a bundle and its associated expenditure) that leads to a violation of the subsistence constraint incurs a very large disutility and so will never be chosen.

We work with AP's data and refer to them for a detailed explanation. Compared to their analysis, we restrict ourselves to the narrower definition of village ("locality") because the larger units of analysis ("municipality") may not be contained in either the treatment or the control group. Also, because we are interested in intertemporal within-village price variation, we estimate separate price schedules for the same village in different waves as opposed to one price schedule (estimated across waves) per village. This necessitates being slightly more permissive about data needs, and we estimate prices for all village-good-wave triples that have 20 or more (as opposed to 75 or more) observations. We follow AP in rejecting data for villages where prices strictly increase with quantity sold and where there is insufficient variation in quantities purchased.

We estimate the price schedule for good i in village v at wave t by applying Ordinary Least Squares to

$$\log(\psi_{vti}(q_{vtih})) = b_{vti0} + b_{vti1} \log(x_{vtih}) + \varepsilon_{vtih}. \quad (27)$$

Here, h indexes households and $\psi_{vti}(q_{vtih}) = \mathbb{E}[p_{vti}(x_{vtih}) | x_{vtih}]$, where $p_{vti}(x_{vtih})$ is the unit price corresponding to quantity x_{vtih} , ε is measurement error, and the expected value is taken over the empirical distribution of reported unit prices corresponding to the same quantity purchased of good i in village-wave (v, t) . This is exactly equation (15) in AP except for being estimated at a less aggregated level.

We test GAPP on untreated households in treated villages (for which we estimate price schedules) with observations in more than one wave and who purchased at least one of the three goods. In our final sample, this leaves us with 2488 households in 177 villages.³⁶

We emphasize that GAPP is not vacuously satisfied on these data. Recall that GAPP cannot be violated when two price systems ψ, ψ' are ranked, in the sense that $\psi(x) \geq \psi'(x)$ for all $x \in \mathbb{R}_+^L$. Of the 20,556 possible combinations of pairs of waves encountered by households in the data, about 4% have this feature, and only 20 out of 2488 households exclusively face such price pairs and therefore satisfy GAPP vacuously. Nonetheless, 83% of households pass the GAPP test. Most

36. For 554 of these households we have two observations, for 840 households we have three, for 934 households we have four, and for 160 households we have five. There are so few with five observations because many households were enrolled into the program in the final wave and thus removed from our sample.

TABLE 1
Fraction of GAPP rationalizable consumers revealed preferring the row wave to the column wave

	10/98	03/99	11/99	11/00	2003
10/98		0.035	0.024	0.006	0
03/99	0.913		0.198	0.052	0.015
11/99	0.936	0.686		0.105	0.034
11/00	0.981	0.914	0.847		0.240
2003	1	0.980	0.927	0.520	

TABLE 2
Lower bound of the compensating variation, with 10/98 as the base

	03/99	11/99	11/00
75th percentile	5.36	7.26	11.65
Median	3.27	4.66	6.98
25th percentile	1.58	2.44	4.12

violations were small in the sense of the rationality index ϑ (defined in Section 3.4) being close to 1: fewer than 1% of households were below 0.9, and fewer than 4% were below .95.

We carried out some illustrative welfare analysis, the results of which are displayed in Tables 1 and 2. Table 1 displays the fractions of GAPP-compliant households that reveal prefer a given wave to another wave. Specifically, each cell in the table corresponds to the fraction of GAPP-rationalizable consumers who reveal prefer (directly or indirectly) the price system in the row wave to the price system in the corresponding column wave.³⁷ Notice that the data indicate a strong tendency to prefer price systems in later waves. For example, 91.3% of households reveal prefer prices in 03/99 to those in 10/98; the same is true even more strongly when 10/98 is compared against later waves.

Table 2 provides scale for this welfare improvement. We calculate, for each household, the lower bound on the compensating variation, with the price system faced by the household at 10/98 as the base.³⁸ These values are then ranked. Since more than 90% of households reveal prefer (price systems at) subsequent waves to 10/98, the lower bound of the compensating variation must be positive for more than 90% of households. For example, between 03/99 and 10/98, the median compensating variation is 3.27; thus, based on its observed behaviour, one could remove 3.27 from this household in 03/99 and still leave it as well off in 03/99 as in 10/98. Note that the values in this table are not small, given that the household median expenditure in 10/98 on the items considered is 27.48.

These results are consistent with AP's finding that the change in the income distribution induced by Progresá caused a change in sellers' intensity of price discrimination. As a result, poorer households faced higher average prices and wealthier households faced lower ones; since Progresá was means-tested, untreated households fall into the latter category. Thus, the general equilibrium effects of the program could be the reason for the welfare improvements observed in untreated households.

37. Note that the (indirect) revealed preference relation \succeq_p^* uses demand information at *all* waves in each binary comparison; see the definition of \succeq_p^* in Section 4.

38. The formula for the lower bound when prices are non-linear is in the Supplementary Appendix A.5.

7.2. RAUM: testing and welfare analysis on household expenditure data

We test the RAUM and conduct welfare analyses on two repeated cross-sectional data sets: the UK Family Expenditure Survey (FES) and the Canadian Surveys of Household Spending (SHS). Our aim is to show that the data support the model and to demonstrate that the estimated welfare bounds are informatively tight.

We first analyse the FES which is widely used in the non-parametric demand estimation literature (for instance, by Blundell, Browning and Crawford (2008), KS, Hoderlein and Stoye (2014), and Adams (2020)). In the FES, about 7000 households are interviewed each year and they report their consumption expenditures in different commodity groups. Following Blundell *et al.* (2008), we derive the real consumption level for each commodity group by deflating it with a price index for that group (which is taken from the annual Retail Prices index). Again following them, we restrict attention to households with cars and children, leaving us with roughly 25% of the original data. We implement tests for 3, 4, and 5 composite goods. Blundell *et al.* (2008) analyse the coarsest partition of three goods—food, services, and non-durables—and we use their replication files. As in KS, we introduce more commodities by first separating out clothing and then alcoholic beverages from the nondurables.

The data are the sample analogue of $\mathcal{D} = \{(p^t, \hat{\pi}^t)\}_{t=1}^T$ (see Section 5.2). We re-iterate the point that, even though this data set is *not* iso-expenditure, we can directly test the RAUM on this data; this contrasts with testing the RUM on this data, which cannot be done directly and must involve a further procedure to estimate an iso-expenditure data set.

We implement the test in blocks of 6 years, i.e., we set $T=6$. We avoid covering a longer period partly due to the computational demands of calculating A (the matrix of GARP-consistent types; see (18)),³⁹ but also because a time-invariant distribution of augmented utility functions is only plausible over shorter time horizons, for example because of long term first-order changes to the UK income distribution (Jenkins, 2016).

Table 3 displays our results: columns correspond to different blocks of 6 years and rows contain the values of the test statistic and the corresponding p -values. The test statistic J_N is defined by (23), with the identity matrix serving as Ω . Note that for the years 90-95, the test statistic is zero; this means that the sample distribution $\hat{\pi}$ satisfies the rationality condition in Theorem 3 exactly. That is, there is a distribution ν on GARP-consistent types such that $\hat{\pi} = A\nu$. Apart from this case, the sample distribution does not exactly satisfy the rationality condition and so the test statistic is strictly positive; nonetheless, the p -values make it very clear that, overall, our model is not rejected by the FES data.

We also estimated the bounds $[\underline{\mathcal{N}}_{p^t \succeq_p^* p^{t'}}, \overline{\mathcal{N}}_{p^t \succeq_p^* p^{t'}}]$ (as defined by (20) and (21)) on the proportion of households that are revealed better off at prices p^t than at prices $p^{t'}$. For brevity, Table 4 presents a few representative estimates using data from 1975 to 1980. The second column are the bounds obtained by calculating $\mathbb{1}_{p^t \succeq_p^* p^{t'}} \nu$ from the (not necessarily unique) values of ν that minimize the test statistic (23). In two cases, this estimate is unique. Applying the procedure for calculating confidence intervals in Section 6.2, we obtain the intervals displayed (which necessarily contain the estimated bounds). We note that the width of these intervals is less than 0.1 throughout, so they are quite informative.⁴⁰

39. That said, new techniques developed in Smeulders, Cherchye and de Rock (2021) have significantly reduced the computational demands of the problem.

40. Note that, the true values of the proportion of the population satisfying $p^t \succ_p^* p^{t'}$ and $p^{t'} \succ_p^* p^t$ typically add up to strictly less than 1 because there is no revealed price preference relation between p^t and $p^{t'}$ for some share of the population. See, for example, type 1 consumers in Example 4.

TABLE 3
Test statistics, p-values for six budget sequences of the FES

		Year blocks									
		75-80	76-81	77-82	78-83	79-84	80-85	81-86	82-87	83-88	84-89
3 Goods	Test statistic (J_N)	0.337	0.917	0.899	0.522	0.018	0.082	0.088	0.095	0.481	0.556
	p -value	0.04	0.34	0.55	0.59	0.99	0.67	0.81	0.91	0.61	0.48
4 Goods	Test statistic (J_N)	0.4	0.698	0.651	0.236	0.056	0.036	0.037	0.043	0.043	0.232
	p -value	0.25	0.58	0.63	0.91	0.96	0.99	0.96	0.95	0.99	0.68
5 Goods	Test statistic (J_N)	0.4	0.687	0.705	0.329	0.003	0.082	0.088	0.104	0.103	0.144
	p -value	0.3	0.66	0.68	0.88	0.999	0.96	0.79	0.85	0.9	0.83

		Year blocks									
		85-90	86-91	87-92	88-93	89-94	90-95	91-96	92-97	93-98	94-99
3 Goods	Test statistic (J_N)	0.027	1.42	2.94	1.51	1.72	0	0.313	0.7	0.676	0.26
	p -value	0.69	0.3	0.18	0.24	0.21	1	0.59	0.48	0.6	0.83
4 Goods	Test statistic (J_N)	0.227	0.025	0.157	0.154	0.004	1.01	0.802	0.872	0.904	0.604
	p -value	0.48	0.96	0.8	0.73	0.97	0.21	0.31	0.57	0.65	0.74
5 Goods	Test statistic (J_N)	0.031	0.019	0.018	0.019	0.023	0.734	0.612	0.643	0.634	0.488
	p -value	0.85	0.98	0.97	0.91	0.83	0.22	0.4	0.72	0.78	0.79

Notes: Bootstrap size is $R = 1000$.

TABLE 4
Estimated bounds and confidence intervals for the proportion of consumers who reveal prefer one price to another one in the FES data

Comparison	Estimated bounds	Confidence interval
$p^{1976} \succ_p^* p^{1977}$	[0.150, 0.155]	[0.123, 0.192]
$p^{1977} \succ_p^* p^{1976}$	{0.803}	[0.777, 0.840]
$p^{1979} \succ_p^* p^{1980}$	[0.517, 0.530]	[0.480, 0.562]
$p^{1980} \succ_p^* p^{1979}$	{0.463}	[0.434, 0.499]

Notes: Data used are for 1975–80. Bootstrap size is $R = 1000$.

For our second empirical application using Canadian data, we use the replication kit of Norris and Pendakur (2013, 2015). Like the FES, the SHS is a publicly available, annual data set of household expenditures in different categories. We study annual expenditures in five categories that constitute a large share of the overall expenditure on nondurables: food purchased (at home and in restaurants), clothing and footwear, health and personal care, recreation, and alcohol and tobacco. As Table 4 shows, the SHS data allow us to analyse the data separately for the nine most populous provinces. The number of households in each province-year range from 291 (Manitoba, 1997) to 2515 (Ontario, 1997). We use province-year prices indices (as constructed by Norris and Pendakur (2015)) and deflate them using province-year CPI data from Statistics Canada to get real price indices.

Table 5 displays the test statistics and associated p -value for each province and every 6-year block. Compared to the FES data, there are two notable differences. The first is that many more test statistics are exactly zero; that is, the observed choice frequencies are rationalized by the RAUM. The second is that, for a small proportion of year blocks, there are statistically significant positive test statistics (in particular, the last three columns for British Columbia). Nonetheless, the p -values taken together do not reject the model if multiple testing is taken into account; for example, step-down procedures would terminate at the first step (i.e. Bonferroni adjustment). Finally, we can also estimate the proportion of the population with a revealed preference for 1 year's prices over another. We provide an illustration in Table 6; notice that the confidence intervals are informative, with a width no greater than 0.15.

TABLE 5
Test statistics and p-values for sequences of six budgets of the SHS. Bootstrap size is R = 1,000.

Provinces		Year blocks							
		97-02	98-03	99-04	00-05	01-06	02-07	03-08	04-09
Alberta	Test statistic (J_N)	0.07	0	0	0	0	0	0.003	4.65
	p-value	0.94	1	1	1	1	1	0.98	0.04
British Columbia	Test statistic (J_N)	0.89	0.56	0.05	0.07	0.05	6.23	8.87	8.71
	p-value	0.47	0.47	0.98	0.96	0.97	0.05	0.02	0.01
Manitoba	Test statistic (J_N)	0	0	0	0	0	0	0.01	0.01
	p-value	1	1	1	1	1	1	1	1
New Brunswick	Test statistic (J_N)	0.08	0.05	0	0	0	0.60	0.58	0.57
	p-value	0.94	0.94	1	1	1	0.58	0.79	0.68
Newfoundland	Test statistic (J_N)	0.10	0.32	0.29	0.29	0.38	3.08	2.30	2.08
	p-value	0.85	0.90	0.91	0.87	0.81	0.21	0.35	0.27
Nova Scotia	Test statistic (J_N)	0.05	0.03	0	0	0	0	0.93	1.02
	p-value	0.97	0.98	1	1	1	1	0.69	0.58
Ontario	Test statistic (J_N)	0.064	0.040	0.035	0	0	0	0	0
	p-value	0.98	0.95	0.91	1	1	1	1	1
Quebec	Test statistic (J_N)	0.11	0	0	0	0	0.51	0.54	0.49
	p-value	0.88	1	1	1	1	0.67	0.67	0.65
Saskatchewan	Test statistic (J_N)	0	0	0	0	0	0.02	0.02	0
	p-value	1	1	1	1	1	1	1	1

TABLE 6
Estimated bounds and confidence intervals for the proportion of consumers who reveal prefer one price to another one in the SHS data. Data used are for 1997–2002 in British Columbia. Bootstrap size is R = 1,000.

Comparison	Estimated bounds	Confidence interval
$p^{1998} \succ_p^* p^{2001}$	{0.099}	[0.073, 0.125]
$p^{2001} \succ_p^* p^{1998}$	{0.901}	[0.875, 0.927]
$p^{1999} \succ_p^* p^{2002}$	[0.299, 0.341]	[0.272, 0.385]
$p^{2002} \succ_p^* p^{1999}$	[0.624, 0.701]	[0.594, 0.728]

8. CONCLUSION

We propose a revealed price preference relation that generates a nonparametric ranking of price vectors; a consistency (no-cycles) condition in this relation characterizes an augmented utility model in which consumers get utility from consumption and disutility from expenditure. This model is a natural generalization of quasilinearity and, furthermore, captures some prominent behavioural models of consumption. The model is also flexible enough to accommodate non-linear prices, discrete choice and other consumption environments. We develop the theoretical basis for welfare analysis in our model.

We generalize our model to a random utility context which is suitable for welfare analysis using repeated cross-sectional (as opposed to single-agent) data and show how to statistically test this RAUM. A strength of this model is that it can be directly taken to household expenditure data in contrast to the standard RUM which requires an additional round of estimation to account for the dependence of expenditure on prices. We develop novel econometric theory to determine the proportion of consumers who are made better or worse off by a price change. This theory—which derives bounds on linear transforms of partially identified vectors—is a standalone contribution which has broader applications beyond those in this article.

Finally, we operationalize both the deterministic and random versions of our model in separate applications to single-agent and repeated cross-sectional data. We confirm that our model is supported by data and can be used for meaningful welfare analysis.

Acknowledgments. We are very grateful for comments provided by the editor, the referees and numerous seminar, conference participants. We also thank Samuel Norris and Krishna Pendakur for generously sharing their data with us and John Hoddinott for patiently explaining Progresa. Deb thanks the Social Sciences and Humanities Research Council (SSHRC) for their continued financial support and the Cowles Foundation for their hospitality. Stoye acknowledges support from National Science Foundation (NSF) grants SES-1260980 and SES-1824375 as well as from Cornell University's Economics Computer Cluster Organization, which was partially funded through NSF Grant SES-0922005. Part of this work was done while Stoye was at Bonn University and the Hausdorff Center for Mathematics. We thank Michael Sullivan and Matthew Thirkettle for excellent research assistance.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online. And the replication packages are available at <https://dx.doi.org/10.5281/zenodo.6564446>.

Data Availability Statement

The data and code underlying this research is available at Zenodo via the following link: <https://doi.org/10.5281/zenodo.6564446>.

REFERENCES

- ADAMS, A. (2020), "Mutually Consistent Revealed Preference Bounds", *American Economic Journal: Microeconomics*, **12**, 42–74.
- AFRIAT, S. N. (1967), "The Construction of Utility Functions from Expenditure Data", *International Economic Review*, **8**, 67–77.
- AFRIAT, S. N. (1973), "On a System of Inequalities in Demand Analysis: An Extension of the Classical Method", *International Economic Review*, **14**, 460–472.
- AMADOR, M., WERNING, I. and ANGELETOS, G.-M. (2006), "Commitment vs. Flexibility", *Econometrica*, **74**, 365–396.
- ATTANASIO, O. and PASTORINO, E. (2020), "Nonlinear Pricing in Village Economies", *Econometrica*, **88**, 207–263.
- BERRY, S., LEVINSOHN, J. and PAKES, A. (1995), "Automobile Prices in Market Equilibrium", *Econometrica*, **63**, 841–890.
- BHATTACHARYA, D. (2015), "Nonparametric Welfare Analysis for Discrete Choice", *Econometrica*, **83**, 617–649.
- BLOW, L., BROWNING, M. and CRAWFORD, I. (2008), "Revealed Preference Analysis of Characteristics Models", *Review of Economic Studies*, **75**, 371–389.
- BLUNDELL, R., BROWNING, M. and CRAWFORD, I. (2008), "Best Nonparametric Bounds on Demand Responses", *Econometrica*, **76**, 1227–1262.
- BORDALO, P., GENNAIOLI, N. and SHLEIFER, A. (2013), "Salience and Consumer Choice", *Journal of Political Economy*, **121**, 803–843.
- CHETTY, R. (2009), "Sufficient Statistics for Welfare Analysis: A Bridge between Structural and Reduced-Form Methods", *Annual Review of Economics*, **1**, 451–488.
- CHETTY, R., LOONEY, A. and KROFT, K. (2009), "Salience and Taxation: Theory and Evidence", *American Economic Review*, **99**, 1145–1177.
- DE PAULA, A., RICHARDS-SHUBIK, S. and TAMER, E. (2018), "Identifying Preferences in Networks With Bounded Degree", *Econometrica*, **86**, 263–288.
- DUPAS, P. and ROBINSON, J. (2013), "Why Don't the Poor Save More? Evidence from Health Savings Experiments", *American Economic Review*, **103**, 1138–1171.
- FELDMAN, N. E. (2010), "Mental Accounting Effects of Income Tax Shifting", *Review of Economics and Statistics*, **92**, 70–86.
- GABAIX, X. (2014), "A Sparsity-Based Model of Bounded Rationality", *Quarterly Journal of Economics*, **129**, 1661–1710.
- (2019), "Behavioral Inattention", in Bernheim, B. D., DellaVigna, S. and Laibson, D. (eds) *Handbook of Behavioral Economics: Applications and Foundations 1*, vol. **2** (North-Holland) 261–343.
- GALICHON, A. and HENRY, M. (2011), "Set Identification in Models with Multiple Equilibria", *Review of Economic Studies*, **78**, 1264–1298.
- HASTINGS, J. and SHAPIRO, J. M. (2018), "How are SNAP Benefits Spent? Evidence from a Retail Panel", *American Economic Review*, **108**, 3493–3540.
- HASTINGS, J. S. and SHAPIRO, J. M. (2013), "Fungibility and Consumer Choice: Evidence from Commodity Price Shocks", *Quarterly Journal of Economics*, **128**, 1449–1498.
- HAUSMAN, J. A. and NEWEY, W. K. (2016), "Individual Heterogeneity and Average Welfare", *Econometrica*, **84**, 1225–1248.
- HEIDHUES, P. and KŐSZEGI, B. (2008), "Competition and Price Variation when Consumers are Loss Averse", *American Economic Review*, **98**, 1245–68.

- HODDINOTT, J., SKOUFIAS, E. and WASHBURN, R. (2000), "The Impact of PROGRESA on Consumption: A Final Report" (Discussion Paper).
- HODERLEIN, S. and STOYE, J. (2014), "Revealed Preferences in a Heterogeneous Population", *The Review of Economics and Statistics*, **96**, 197–213.
- JENKINS, S. P. (2016), "The Income Distribution in the UK", in Dean, H. and Platt, L. (eds) *Social Advantage and Disadvantage* (Oxford, UK: Oxford University Press) 135–160.
- KITAMURA, Y. and STOYE, J. (2018), "Nonparametric Analysis of Random Utility Models", *Econometrica*, **86**, 1883–1909.
- (2019), "Nonparametric Counterfactuals of Random Utility Models" (Discussion Paper, Cornell University) arXiv:1902.08350.
- KLEVEN, H. (2021), "Sufficient Statistics Revisited", *Annual Review of Economics*, **13**, 515–538.
- KÓSZEGI, B. and RABIN, M. (2006), "A Model of Reference-Dependent Preferences", *Quarterly Journal of Economics*, **121**, 1133–1165.
- LAZZATI, N., QUAH, J. K.-H. and SHIRAI, K. (2018), "Nonparametric Analysis of Monotone Choice" (Working Paper).
- MCFADDEN, D. and RICHTER, K. (1991), "Stochastic Rationality and Revealed Stochastic Preference", in Chipman, J., McFadden, D. and Richter, K. (eds) *Preferences, Uncertainty and Rationality* (Boulder: Westview Press) 161–186.
- MCFADDEN, D. L. (2005), "Revealed Stochastic Preference: A Synthesis", *Economic Theory*, **26**, 245–264.
- NEVO, A. (2000), "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand", *Journal of Economics and Management Strategy*, **9**, 513–548.
- NORRIS, S. and PENDAKUR, K. (2013), "Imputing Rent in Consumption Measures, with an Application to Consumption Poverty in Canada, 1997–2009", *Canadian Journal of Economics*, **46**, 1537–1570.
- (2015), "Consumption Inequality in Canada, 1997 to 2009", *Canadian Journal of Economics*, **48**, 773–792.
- PRELEC, D. and LOEWENSTEIN, G. (1998), "The Red and the Black: Mental Accounting of Savings and Debt", *Marketing Science*, **17**, 4–28.
- SAKAI, Y. (1977), "Revealed Favorability, Indirect Utility, and Direct Utility", *Journal of Economic Theory*, **14**, 113–129.
- SAKOVICS, J. (2011), "Reference Distorted Prices", *Quantitative Marketing and Economics*, **9**, 339–363.
- SMEULDERS, N., CHERCHYE, L. and DE ROCK, B. (2021), "Nonparametric Analysis of Random Utility Models: Computational Tools for Statistical Testing", *Econometrica*, **89**, 437–455.
- THALER, R. H. (1999), "Mental Accounting Matters", *Journal of Behavioral Decision Making*, **12**, 183–206.
- VARIAN, H. R. (1983), "Non-parametric Tests of Consumer Behaviour", *Review of Economic Studies*, **50**, 99–110.