

Goods versus characteristics: dimension reduction and revealed preference

IFS Working Paper W12/02

Matthew Polisson

Goods Versus Characteristics: Dimension Reduction and Revealed Preference*

Matthew Polisson[†]

March 1, 2012

Abstract

This paper compares the goods and characteristics models of the consumer within a non-parametric revealed preference framework. Of primary interest is to make a comparison on the basis of predictive success that takes into account dimension reduction. This allows us to nonparametrically identify the model which best fits the data. We implement these procedures on household panel data from the UK milk market. The primary result is that the better fit of the characteristics model is entirely attributable to dimension reduction.

JEL Classification: D11, D12.

Keywords: Characteristics, demand, dimension reduction, nested models, revealed preference.

1 Introduction

This paper compares two models that are central to empirical work in applied microeconomics, namely the goods model of the consumer (Deaton and Muellbauer (1980); Banks, Blundell, and Lewbel (1997)) and the characteristics model of the consumer (McFadden (1973); Berry, Levinsohn, and Pakes (1995)). The former posits that an individual has preferences over goods, and the latter that an individual has preferences over characteristics, giving rise to derived preferences over goods (see Gorman (1956) and Lancaster (1966) for early descriptions of the characteristics model). Both models are parsimonious and empirically falsifiable, but the characteristics model has a further advantage in that it reduces the dimensions of the consumer choice problem.

Since the characteristics model is a special case of the goods model, it necessarily performs weakly worse from a *consistency* standpoint. It is obvious that additional theoretical structure begets additional empirical restrictions, which then implies that a general model is at least as likely as a special case to be consistent with the data. While this raises a broader question about when to add structure to a model more generally, in this paper we focus exclusively on the goods and characteristics models in order to explore the effects of dimension reduction. We make a comparison by assessing consistency in light of the criteria for consistency taking into account dimension reduction. In doing so, we identify the value added by the characteristics structure in explaining consumer behaviour.

The conventional econometric approach has been to first impose functional form and heterogeneity restrictions, and then to estimate both models, comparing fit. However, even with flexible functional forms and allowing for heterogeneity, this approach necessarily conflates primitive and ancillary hypotheses. Instead, we proceed empirically in the tradition of nonparametric revealed

*For helpful comments and useful suggestions, I am grateful to Martin Browning, Ian Crawford, Rachel Griffith, Valérie Lechene, and John Quah, to seminar participants at the IFS, Leicester, and Oxford, and to conference delegates at the 2011 SAET Conference. I am also grateful to Kantar and the IFS for providing access to the data. Funding from the ESRC/MRC/NIHR, Reference Number G0802297, is gratefully acknowledged. A previous draft of this paper was circulated under the title “Goods Versus Characteristics: Revealed Preference Procedures for Nested Models”.

[†]Department of Economics, University of Leicester, University Road, Leicester, LE1 7RH and Institute for Fiscal Studies, 7 Ridgmount Street, London, WC1E 7AE. E-mail: matthew.polisson@le.ac.uk.

preference, for which there is a well-established methodology of falsification and a vast supporting literature (Afriat (1967); Diewert (1973); Varian (1982)). Necessary and sufficient conditions are well-defined and easy to implement using linear programming techniques for both the goods and characteristics models (see Blow, Browning, and Crawford (2008) for the latter). These nonparametric tests exhaust the pure empirical implications of each theory and avoid the identification problems associated with conventional parametric approaches.

A drawback of nonparametric revealed preference procedures is that they very often fail to give rise to stringent, or even non-trivial, empirical restrictions. Without sufficient price relative to expenditure variation, the ability to falsify is minimal. In a non-statistical sense, nonparametric revealed preference conditions very often lack restrictive (and predictive) *power*. This problem is well-known, and several treatments have been advanced, most notably by Bronars (1987). However, in this paper we are able to exploit the power differential between the goods model and the weakly more restrictive characteristics model. We make a comparison on the basis of predictive success, which requires that we relate set predictions across models. To do this, we make use of an axiomatic approach developed by Selten (1991) and already extended to the revealed preference framework by Beatty and Crawford (2009). We relate consistency and power in order to establish *fit*, again in a non-statistical sense, and then make a comparison. Our primary innovation is to account for dimension reduction when making this comparison.

We implement these procedures on household panel data from the UK milk market. This highlights a relevant nutritional application, in which we aim to identify the relative importance of nutrients in explaining food consumption. There are an abundance of alternative markets, e.g., housing, automobile, education, and labour, but we focus on milk to simplify the implementation and to establish continuity with previous work in this area. Our data are from the Kantar UK Worldpanel, which is comprised of scanner data from a representative rolling household panel.

Our results suggest that on average the characteristics model outperforms the goods model in terms of fit. However, the better fit of the characteristics model is entirely attributable to its dimension-reducing properties, which suggests that other dimension-reducing structures, e.g., weak separability or alternative technological mappings from goods to characteristics, may have performed at least as well. Although there is a great deal of heterogeneity in our results, we find that consistency, power, and fit are largely unexplained by observable household attributes.

The organisation of this paper is as follows. Section 2 develops the basic theory. We first situate the analysis within a traditional demand framework and describe the critical features and empirical implications of the goods and characteristics models. We then outline the revealed preference conditions for these models, as well as a power concept already developed for the goods model. Next, we extend this concept to the characteristics model, and in doing so, we discuss our priors and how to handle dimension reduction. We then briefly review the axiomatisation of predictive success. Section 3 outlines the empirical implementation. Here we introduce the nutritional application and describe the data. Section 4 presents the results. Section 5 concludes.

2 Theory

2.1 Framework

The primal problem in the goods model is to

$$\max_{\mathbf{q}} u(\mathbf{q}) \text{ subject to } \mathbf{p}'\mathbf{q} \leq x, \quad (1)$$

where $\mathbf{q} \in \mathbb{R}_+^K$ are consumption goods, $\mathbf{p} \in \mathbb{R}_{++}^K$ are prices, and $x \in \mathbb{R}$ is income, and where the utility function $u(\cdot) : \mathbb{R}_+^K \rightarrow \mathbb{R}$ is differentiable, non-satiated, and concave in \mathbf{q} . The primal problem in the characteristics model is to

$$\max_{\mathbf{q}} v(\mathbf{z}) \text{ subject to } \mathbf{z} = F(\mathbf{q}) \text{ and } \mathbf{p}'\mathbf{q} \leq x, \quad (2)$$

where $\mathbf{z} \in \mathbb{R}_+^J$ are characteristics, and where the utility function $v(\cdot) : \mathbb{R}_+^J \rightarrow \mathbb{R}$ is differentiable, non-satiated, and concave in \mathbf{z} , and the mapping from goods to characteristics $F(\cdot) : \mathbb{R}_+^K \rightarrow \mathbb{R}_+^J$ is

differentiable, nondecreasing, and concave in \mathbf{q} . We adopt a standard linear mapping $F(\mathbf{q}) = \mathbf{A}'\mathbf{q}$, where $\mathbf{A} \in \mathbb{R}_+^{K \times J}$, $J < K$, and \mathbf{A} has full rank by construction.¹ We assume a linear form but note that what follows applies to more general technologies.

The solutions to (1) and (2) are the demand systems $\mathbf{q}(\mathbf{p}, x) : \mathbb{R}_{++}^K \times \mathbb{R}_{++} \rightarrow \mathbb{R}_+^K$ and $\tilde{\mathbf{q}}(\mathbf{p}, x; \mathbf{A}) : \mathbb{R}_{++}^K \times \mathbb{R}_{++} \rightarrow \mathbb{R}_+^K$, respectively, which are both differentiable and zero homogenous in (\mathbf{p}, x) , satisfy adding-up, and give rise to a Slutsky matrix that is symmetric and negative semi-definite. However, there is a further empirical restriction on $\tilde{\mathbf{q}}(\cdot)$ that the number of goods consumed cannot exceed the number of characteristics, unless characteristics budget sets are uniquely aligned.

This is best illustrated diagrammatically in Figure 1 using an example with three goods and two characteristics. Each ray emanating from the origin has a gradient corresponding to the ratio of characteristic 2 to characteristic 1 embodied by a particular good. A consumer is therefore constrained by prices, income, and technology. In the example in Figure 1a, the characteristics budget set restricts the individual to consuming only good 1, only good 2, only good 3, some of good 1 and some of good 2, or some of good 2 and some of good 3. Consuming some of good 1 and some of good 3 and consuming all goods are both ruled out in this particular scenario. One can imagine other characteristics budget sets that give rise to similar dimension-reducing restrictions. In the example in Figure 1b, the individual is free to consume in any combination. These restrictions are neatly encapsulated by the rank condition

$$\text{rank}(\mathbf{A}^+ \sim \mathbf{p}^+) = \text{rank}(\mathbf{A}^+), \quad (3)$$

where \mathbf{p}^+ is a sub-vector of \mathbf{p} and \mathbf{A}^+ is a sub-matrix of \mathbf{A} for which the corresponding elements of \mathbf{q} are zero, and where \sim denotes the horizontal concatenation of \mathbf{A}^+ and \mathbf{p}^+ . Since $J < K$ by construction, the rank condition normally restricts consumption to fewer than the full set of goods.

Preferences over characteristics, combinations of which are embodied by goods, imply derived preferences over goods, i.e., there exists a derived utility function $\tilde{u}(\cdot) : \mathbb{R}_+^K \rightarrow \mathbb{R}$ such that $\tilde{u}(\mathbf{q}) = v(\mathbf{A}'\mathbf{q})$, which means that the characteristics model is a special case of the goods model. In this way, the goods model nests the characteristics model, and as consequence, any observable empirical restrictions are also nested. We exploit this nested structure to identify which model better fits the data.

2.2 Revealed preference conditions

2.2.1 Goods model

Since we observe prices and quantities, and hence expenditures, the restrictions outlined in the previous section are, in principle, empirically falsifiable. The conventional econometric approach has been to first impose functional form and heterogeneity restrictions. However, parametric techniques necessarily conflate primitive and ancillary hypotheses. Instead, we proceed empirically in a nonparametric revealed preference sense, for which there is a well-established methodology of falsification and a vast supporting literature (see Samuelson (1938, 1948), Houthakker (1950), and Richter (1966) for original formulations of the argument, Afriat (1967) and Diewert (1973) for necessary and sufficient conditions under which a set of observations is consistent with the standard goods model of the consumer, and Varian (1982, 1983) for operationalisable and computationally feasible tests).

Let $\mathbf{q}_t \in \mathbb{R}_+^K$ be observed goods bought at prices $\mathbf{p}_t \in \mathbb{R}_{++}^K$ in period $t \in \{1, \dots, T\}$. If the set of observations $(\mathbf{p}_t, \mathbf{q}_t)$, $t = 1, \dots, T$, is consistent with the goods model of the consumer, it can be rationalised as follows:

Definition 1 A utility function $u(\cdot) : \mathbb{R}_+^K \rightarrow \mathbb{R}$ q -rationalises the set of observations $(\mathbf{p}_t, \mathbf{q}_t)$, $t = 1, \dots, T$, if $u(\mathbf{q}_t) \geq u(\mathbf{q})$ for all \mathbf{q} such that $\mathbf{p}'_t \mathbf{q}_t \geq \mathbf{p}'_t \mathbf{q}$.

¹Lancaster (1966) discusses three alternative technological structures for the linear characteristics model: (1) when $J = K$, the characteristics model collapses to the goods model with transformed prices; (2) when $J > K$, a \mathbf{q} may not exist for every \mathbf{z} ; and (3) when $J < K$, a \mathbf{q} exists for every \mathbf{z} , and an infinite number of \mathbf{q} may exist for some \mathbf{z} . This third non-trivial scenario is most economically interesting since it reduces the dimensions of the choice problem for the consumer.

The necessary and sufficient conditions for q -rationalisation are related by Afriat's theorem as follows:

Theorem 1 *The following statements are equivalent:*

1. *There exists a utility function $u(\cdot) : \mathbb{R}_+^K \rightarrow \mathbb{R}$ that is continuous, increasing, and concave which q -rationalises the set of observations $(\mathbf{p}_t, \mathbf{q}_t)$, $t = 1, \dots, T$.*
2. *There exist $U_t \in \mathbb{R}$, $\lambda_t \in \mathbb{R}_{++}$, $t = 1, \dots, T$,*

$$U_s \leq U_t + \lambda_t \mathbf{p}'_t (\mathbf{q}_s - \mathbf{q}_t) \quad \forall s, t = 1, \dots, T. \quad (4)$$

Proof. See Afriat (1967), Varian (1982, 1983), and Foster, Scarf, and Todd (2004). ■

The main implication of the theorem is that if there exists a solution to the system of T^2 linear inequalities and $2T$ unknowns in (4), then the set of observations is q -rationalisable, or consistent with the goods model of the consumer.²

2.2.2 Characteristics model

Let $\mathbf{z}_t \in \mathbb{R}_+^J$ be observed characteristics in period $t \in \{1, \dots, T\}$, and let $\mathbf{A} \in \mathbb{R}_+^{K \times J}$ be a linear technology mapping from \mathbf{q}_t to \mathbf{z}_t according to $\mathbf{z}_t = \mathbf{A}'\mathbf{q}_t \quad \forall t = 1, \dots, T$. If the set of observations $(\mathbf{p}_t, \mathbf{q}_t)$, $t = 1, \dots, T$, and the technology \mathbf{A} are consistent with the linear characteristics model of the consumer, they can be rationalised as follows:

Definition 2 *A utility function $v(\cdot) : \mathbb{R}_+^J \rightarrow \mathbb{R}$ z -rationalises the set of observations $(\mathbf{p}_t, \mathbf{q}_t)$, $t = 1, \dots, T$, for the technology \mathbf{A} , if $v(\mathbf{z}_t) = v(\mathbf{A}'\mathbf{q}_t) \geq v(\mathbf{A}'\mathbf{q}) = v(\mathbf{z})$ for all \mathbf{q} such that $\mathbf{p}'_t \mathbf{q}_t \geq \mathbf{p}'_t \mathbf{q}$.*

The necessary and sufficient conditions for z -rationalisation are related as follows:

Theorem 2 *The following statements are equivalent:*

1. *There exists a utility function $v(\cdot) : \mathbb{R}_+^J \rightarrow \mathbb{R}$ that is continuous, non-satiated, and concave which z -rationalises the set of observations $(\mathbf{p}_t, \mathbf{q}_t)$, $t = 1, \dots, T$, for the technology \mathbf{A} .*
2. *There exist $V_t \in \mathbb{R}$, $\lambda_t \in \mathbb{R}_{++}$, $\pi_t \in \mathbb{R}^J$, $t = 1, \dots, T$, such that*

$$V_s \leq V_t + \pi'_t (\mathbf{A}'\mathbf{q}_s - \mathbf{A}'\mathbf{q}_t) \quad \forall s, t = 1, \dots, T, \quad \text{and} \quad (5)$$

$$\mathbf{A}\pi_t \leq \lambda_t \mathbf{p}_t \quad \forall t = 1, \dots, T, \quad (6)$$

with (6) binding for any $q_{kt} > 0$.

Proof. See Blow, Browning, and Crawford (2008). ■

The main implication of the theorem is that if there exists a solution to the $T(K + T)$ linear (in)equalities and $T(J + 2)$ unknowns in (5) and (6), then the set of observations is z -rationalisable, or consistent with the linear characteristics model of the consumer for a given technology.³

So far we have assumed that all prices are observed. In empirical practice, it is often the case that we only observe the prices of goods purchased, and then we impute the prices of goods not purchased. Blow, Browning, and Crawford (2008) develop a set of revealed preference conditions for the characteristics model that allow for missing prices by removing the inequality restrictions on the prices of goods not purchased in (6). The conditions are necessarily weaker since the prices of goods not purchased can be set arbitrarily high. Blow, Browning, and Crawford (2008) further note that it is possible to construct virtual prices, where individuals are just on the verge of buying, and also that setting the technology \mathbf{A} equal to the identity matrix gives rise to the revealed preference conditions for the goods model with missing prices. In what follows, we assume that all prices are observed until explicitly noted otherwise.

²While linear programming techniques are computationally viable for small-dimensional problems, they become increasingly cumbersome for large-dimensional problems. An alternative strategy is to test for violations of the Generalised Axiom of Revealed Preference (GARP), which is computationally very rapid. Further note that the full theorem implies that non-satiation versus continuity, monotonicity, and concavity are indistinguishable within a finite set of observations.

³Note as before that the full theorem implies that non-satiation versus continuity, non-satiation, and concavity are indistinguishable within a finite set of observations.

2.3 Power of revealed preference

Since the goods model nests the characteristics model, z -rationalisation implies q -rationalisation.⁴ As a result, aggregating over a heterogeneous sample, the pass rate for the goods model is always at least as high as the pass rate for the characteristics model. In other words, the characteristics model necessarily performs weakly worse than the goods model from a consistency standpoint. A challenge motivating this paper has been to formally identify the value added by the characteristics structure in light of these pass rates. In order to do this, we explore power differentials in the revealed preference conditions for the goods and characteristics models, which highlights the importance of dimension reduction. But first we consider the power of revealed preference more generally.

The necessary and sufficient revealed preference conditions outlined in the previous section are exhaustive, meaning that no further derivable empirical implications exist unless we impose additional structure. In choosing not to do so, we avoid conflating hypotheses, and we allow for maximal heterogeneity, both of which are desirable. However, the power of these tests depends vitally upon variations in the data. More specifically, we require sufficient price relative to expenditure variation such that implied budget hyperplanes intersect, which has the potential to generate cycles, and hence the power to falsify. This problem has long been noted, and several treatments have been advanced (see Andreoni and Harbaugh (2008) for a review). Nonetheless, we are able to exploit structure and power differentials between the goods and characteristics models in order to establish comparative fit. The method we adopt is similar in spirit to that of Bronars (1987), but very closely follows upon Beatty and Crawford (2009), who have recently developed an axiomatic approach based on Selten (1991).

To illustrate, we begin with a simple two-good, two-period numerical example. Let $\mathbf{q}_1 = (0, 14)'$ at $\mathbf{p}_1 = (7, 6)'$ and $\mathbf{q}_2 = (14, 0)'$ at $\mathbf{p}_2 = (6, 7)'$. Figure 2 depicts this set of observations. The solid line represents the implied budget line in period 1, and the dashed line the implied budget line in period 2. The set of observations is q -rationalisable, and since the implied budget lines intersect, this is non-trivial. In fact, at the same prices and implied expenditures, we can define a set of outcomes that are not q -rationalisable according to

$$(\mathbf{q}_1, \mathbf{q}_2) : q_{11} \in (84/13, 12], \quad q_{12} \in [0, 84/13), \quad (7, 6)'\mathbf{q}_1 = 84, \quad (6, 7)'\mathbf{q}_2 = 84. \quad (7)$$

Translating into budget shares gives

$$(w_{11}, w_{12}) : w_{11} \in (7/13, 1], \quad w_{12} \in [0, 6/13), \quad (8)$$

where $w_{11}, w_{12} \in [0, 1]$ are budget shares for good 1 in periods 1 and 2, respectively. The benefit of using budget shares is that with only two goods and two periods, the set of q -rationalisable outcomes, implied by adding-up, can be illustrated diagrammatically as in Figure 3. Following Beatty and Crawford (2009), we define a relative *area* $a \in [0, 1]$ as the size of the set of q -rationalisable outcomes relative to the size of the set of all feasible outcomes. Loosely speaking, an area is a theory-consistent acceptance region, or target, that lies within the feasible outcome space. In this example, $a = 133/169 \approx 0.787$, which means that about 21.3% of feasible outcomes are ruled out by the revealed preference restrictions.

It is worth being precise about what an area captures. In the data, we only observe prices and quantities, which are all that we require to establish q -rationalisability. However, in order to calculate an area, we must first define a set of feasible outcomes. At the observed prices in a given period, we can define a set of bundles just as affordable as the observed bundle. In other words, the set of feasible bundles in a given period includes all points on the budget hyperplane implied by observed prices and quantities in that period. The feasible outcome space is simply the cartesian product of these sets across periods, i.e.,

$$\mathcal{F} = \times_{t=1}^T \mathcal{F}_t = \times_{t=1}^T \{\mathbf{q} \in \mathbb{R}_+^K : \mathbf{p}'_t \mathbf{q} = \mathbf{p}'_t \mathbf{q}_t\}, \quad (9)$$

where \mathcal{F} denotes the feasible outcome space and \mathcal{F}_t denotes the set of feasible bundles in period t . Beatty and Crawford (2009) therefore note that area power is vitally dependent upon any priors

⁴It is clear that (5) and (6) imply (4). For completeness, note that (a) z -rationalisation with imputed prices implies q -rationalisation with imputed prices; (b) z -rationalisation with missing prices implies q -rationalisation with missing prices; (c) z -rationalisation with imputed prices implies z -rationalisation with missing prices; and (d) q -rationalisation with imputed prices implies q -rationalisation with missing prices.

about the feasible outcome space. For example, the set of feasible outcomes as just described may contain outcomes that are behaviourally irrelevant, and we may want to truncate the feasible outcome space accordingly. The area treatment is therefore not so easily applied to the characteristics model, and indeed whenever the dimensions of consumer choice have been reduced. We discuss this at length in the following section.

Lastly, note that in empirical practice we use Monte Carlo simulations to numerically estimate areas. In the two-good, two-period example above, we are able to construct an area geometrically and solve for it analytically. We use this simple example to develop intuition. However, with additional goods and periods, or with additional complicating structure, we use numerical methods. An analogous numerical procedure draws outcomes uniformly from implied budget hyperplanes and tests for q -rationalisation until convergence is reached. While we require a probabilistic assumption to implement this procedure, what we are trying to estimate is the relative size of a *set* of theory-consistent outcomes. As Beatty and Crawford (2009) explicitly note, area power is not, in concept, a probability measure, though it has the requisite properties. When given this interpretation, it resembles statistical power (see Bronars (1987)),⁵ which has been criticised for considering, as the hypothesis alternative to the goods model of the consumer, that all outcomes on the implied budget hyperplanes are events occurring with equal probabilities. Such criticism is a reminder that an area critically depends upon the set of feasible outcomes against which the model of interest is being compared.

2.4 Power and dimension reduction

Here we discuss our priors over the set of feasible outcomes for the goods and characteristics models and how to accommodate dimension reduction. First consider

$$\mathcal{F} = \times_{t=1}^T \mathcal{F}_t = \times_{t=1}^T \{\mathbf{q} \in \mathbb{R}_+^K : \mathbf{p}'_t \mathbf{q} = \mathbf{p}'_t \mathbf{q}_t\}. \quad (\text{P1})$$

This corresponds to the set of feasible outcomes outlined in the previous section.

Consider a simple three-good, two-characteristic, two-period numerical example. Let $\mathbf{q}_1 = (0, 0, 14)'$ at $\mathbf{p}_1 = (7, 4, 6)'$ and $\mathbf{q}_2 = (14, 0, 0)'$ at $\mathbf{p}_2 = (6, 4, 7)'$. The technology mapping from goods to characteristics is represented by

$$\mathbf{A} = \begin{pmatrix} 1 & 2 \\ 1 & 1 \\ 2 & 1 \end{pmatrix}. \quad (10)$$

Figure 4 depicts this set of observations in characteristics space. Once again, the solid line represents the implied budget line in period 1, and the dashed line the implied budget line in period 2. The set of observations is q -rationalisable, and since the implied budget planes intersect, this is again non-trivial. Under (P1), $a^q \approx 0.816$, where a^q denotes the area of the goods model, which means that about 18.4% of feasible outcomes are ruled out by corresponding revealed preference restrictions.

Furthermore, the set of observations is also z -rationalisable, non-trivially satisfying the revealed preference conditions for the characteristics model. Under (P1), $a^z \approx 0.000$, where a^z denotes the area of the characteristics model, which means that about 100.0% of feasible outcomes are ruled out by the corresponding revealed preference restrictions. It is easy to see why this happens. An implication of the characteristics model is that the number of goods consumed cannot exceed the number of characteristics, unless characteristics budget sets are uniquely aligned. While the sets of feasible bundles in periods 1 and 2 include all points on the corresponding budget planes implied by observed prices and quantities, the set of z -rationalisable bundles is restricted to points on the edges of these planes.⁶ In general, in a model with fewer characteristics than goods, dimension reduction has the effect of increasing the power by an infinite amount,⁷ given a feasible outcome space corresponding to (P1).

⁵In fact, an area is equal to one minus the statistical power measure of Bronars (1987).

⁶In this example, consuming some of good 1 and some of good 3 and consuming all goods are ruled out in both periods.

⁷Increases in power correspond to smaller values of a^i for model $i \in \{g, z\}$ and vice versa.

Losing a dimension of choice should be taken into account when making power comparisons between the goods and characteristics models. We next consider

$$\mathcal{F} = \times_{t=1}^T \mathcal{F}_t = \times_{t=1}^T \{\mathbf{q} \in \mathbb{R}_+^K : \mathbf{p}'_t \mathbf{q} = \mathbf{p}'_t \mathbf{q}_t, \mathbf{p}_t^0 = \infty\}, \quad (\text{P2})$$

where \mathbf{p}_t^0 is a sub-vector of \mathbf{p}_t for which the corresponding elements of \mathbf{q}_t are zero. This is a natural set of feasible outcomes when prices are missing. When unobserved prices are set arbitrarily high, we essentially are able to “control for” dimension reduction. Here we avoid assigning any weight to the increase in the power of the characteristics model achieved purely by reducing the number of dimensions of choice. Under (P2), $a^q = a^z = 1.000$ in the preceding example, which means that about 0.0% of feasible outcomes are ruled out by the revealed preference conditions for both the goods and characteristics models with missing prices. This happens because the feasible outcome space has been reduced to a single dimension, namely observed prices and quantities of goods purchased. While this example is extreme, missing prices give rise to decreases in power in general.⁸

An alternative strategy altogether is to discretise the feasible outcome space. Under (P1), dimension reduction has the effect of increasing the power of the characteristics model by an infinite amount only because the feasible outcome space is continuous, which is of course unrealistic in empirical practice. A solution is to assume that the feasible outcome space is discrete, which better reflects the practical nature of the choice problem and allows us to calculate exact areas.⁹ Consider

$$\mathcal{F} = \times_{t=1}^T \mathcal{F}_t = \times_{t=1}^T \{\mathbf{q} \in \mathbb{Z}_+^K : \mathbf{p}'_t \mathbf{q} = \mathbf{p}'_t \mathbf{q}_t\}.^{10} \quad (\text{P3})$$

Note that in the limit areas calculated under (P3) tend towards those estimated under (P1). Under (P3), in the previous example there are 28 feasible bundles in each period if we restrict consumption to whole units, which gives 784 feasible outcomes, of which 619 are q -rationalisable, such that $a^q \approx 0.790$. Of these 784 feasible outcomes, only 106 are z -rationalisable, such that $a^z \approx 0.135$. If we further restrict consumption to half units in the preceding example, $a^q \approx 0.787$ and $a^z \approx 0.042$, tending towards the (P1) results. As we have shown in this section, the area of a model heavily depends upon the feasible outcome space over which that model is defined. Once we have specified our feasibility priors, the remaining step in identifying the model which best fits the data is to appeal to an axiomatic claim about predictive success.

2.5 Predictive success

After defining an area as a set restriction, Beatty and Crawford (2009) make use of a theorem from Selten (1991), who has put forward a measure of predictive success for any theory that predicts a subset of all possible outcomes. This measure was developed initially for experimental game theory. For our purposes, the relative size of the predicted subset is a^i for model $i \in \{q, z\}$, and a set of observations is either an element of the predicted subset or it is not, which is captured by the pass/fail indicator $r^i \in \{0, 1\}$. Loosely speaking, a^i represents the size of the target, and r^i a hit or a miss of the target. In order to assess the performance of an area theory, we require a predictive success measure $m(r^i, a^i) : \{0, 1\} \times [0, 1] \rightarrow \mathbb{R}$. That $m(\cdot)$ takes as its arguments (r^i, a^i) is relatively uncontroversial, but the functional form for $m(\cdot)$ is entirely debatable. Selten (1991) suggests several desirable properties.

Definition 3 *A predictive success measure $m(r^i, a^i)$ satisfies*

1. *Monotonicity* if $m(1, 0) > m(0, 1)$,
2. *Equivalence* if $m(0, 0) = m(1, 1)$, and

⁸For the characteristics model, rather than setting unobserved prices arbitrarily high, it is sometimes possible to construct virtual prices that just allow for observed behaviour. This necessarily increases the power of the characteristics model.

⁹The revealed preference conditions for the goods and characteristics models do not depend upon assuming a continuous outcome space. As long as discrete consumption is weakly separable from continuous consumption, the familiar results hold (see Polisson and Quah (2011) for a formal treatment of discreteness, separability, and revealed preference).

¹⁰Here we have restricted consumption to whole units, but discreteness can be defined more generally.

3. *Aggregability* if $m(\lambda r_1^i + (1 - \lambda)r_2^i, \lambda a_1^i + (1 - \lambda)a_2^i) = \lambda m(r_1^i, a_1^i) + (1 - \lambda)m(r_2^i, a_2^i)$ for any $\lambda \in [0, 1]$.

Monotonicity states that passing infinitely stringent restrictions is a greater success than failing infinitely lenient restrictions; equivalence states that failing infinitely stringent restrictions and passing infinitely lenient restrictions are equally informative; and aggregability states that a predictive success measure should allow for additive decomposition such that it can be applied easily to a heterogenous sample. Given these axioms, Selten (1991) derives the following result.

Theorem 3 *The predictive success measure $m(r^i, a^i) = r^i - a^i \in [-1, 1]$ satisfies monotonicity, equivalence, and aggregability. If another predictive success measure $\tilde{m}(r^i, a^i)$ also satisfies these axioms, then there exist $\alpha^i \in \mathbb{R}$ and $\beta^i \in \mathbb{R}_{++}$ such that $\tilde{m}(r^i, a^i) = \alpha^i + \beta^i m(r^i, a^i)$.*

Proof. See Selten (1991) and Beatty and Crawford (2009). ■

The main implication of the theorem is not only that $m(r^i, a^i) = r^i - a^i$ satisfies three desirable axioms, but also that any other measure satisfying these axioms is an affine transformation of $r^i - a^i$. This result suggests that $m(r^i, a^i) = r^i - a^i$ is a suitable measure for establishing predictive success. The higher this measure, the more successful the model, and vice versa. When $r^i - a^i$ is close to zero, the set of observations is either passing lenient restrictions or failing stringent restrictions.

From Theorem 3, we know that if $m(r^z, a^z) > m(r^q, a^q)$, then the characteristics model outperforms the goods model, and vice versa. If $m(r^z, a^z) = m(r^q, a^q)$, then the two models are equally predictively successful. This suggests that differencing the two measures is informative about *relative* predictive success. Selten (1991) derives an ordinal counterpart to the cardinal result in Theorem 3, which depends upon an axiom stating that any two theories predicting subsets of all possible outcomes should be compared on the basis of changes in accuracy and precision. Within our framework, $\Delta r = r^z - r^q$ denotes the change in accuracy and $\Delta a = a^z - a^q$ the change in precision. Under these assumptions, it can be shown that $\delta(\Delta r, \Delta a) = \Delta r - \Delta a = m(r^z, a^z) - m(r^q, a^q)$ is a suitable measure for establishing relative predictive success. The higher this measure, the more relatively successful the characteristics model, and vice versa.

3 Implementation

3.1 Application

We implement the procedures outlined in previous sections on household panel data from the UK milk market.¹¹ We choose milk to establish continuity with an existing literature (Blow, Browning, and Crawford (2008); Griffith and Nesheim (2010)) and to highlight a relevant nutritional application. The procedures developed in this paper allow us to identify the relative importance of nutrients (characteristics) in explaining food (goods) consumption.¹² We envisage a model in which individuals trade off between taste and health, which is of particular interest for goods that we may perceive to ultimately become “bads”. The relevant policy interest may be whether an effective tax on fat, for example, manifested as a tax on particular foods, causes people to eat more healthily.

3.2 Data

Our data are from the Kantar UK Worldpanel, which is a representative rolling panel of more than 15,000 households since 2001 and more than 25,000 households since 2006. It contains standard demographic information for each household and its constituent members. It also contains detailed expenditure and product information for fast-moving consumer goods since 2001, including nutritional information since 2006. Each household is equipped with a barcode scanner, which records the precise details associated with each purchase. These data are then linked to manufacturer product databases. The relative strengths of the Kantar data (i.e., in comparison with traditional sources like the Expenditure and Food Survey (EFS), formerly the Family Expenditure Survey

¹¹Programs are available from the author upon request.

¹²Since food contains some nutrients that are healthy in small amounts and unhealthy in large amounts, the characteristics model is a very plausible structure for food and nutrients.

(FES)) are that they follow households over time and that they provide a high level of detail. The relative weaknesses come primarily as a consequence of attrition and non-response. See Griffith and O’Connell (2009) and Leicester and Oldfield (2009) for further descriptions of the Kantar data.

We follow a sample of UK households for a full year from November 1, 2006 to October 31, 2007. We focus on the three main types of conventional non-organic milk: whole, semi-skimmed, and skimmed.¹³ Following Blow, Browning, and Crawford (2008), we aggregate household milk purchases to a monthly level in order to reduce the computational burden and to respect the intertemporal separability implicit in the models we are testing.¹⁴ For each household and month, we observe quantities and expenditures for each type of milk that is purchased, from which we construct unit prices. For types that are not purchased, we impute missing unit prices by taking the medians of observed unit prices by region and month.¹⁵ We round monthly quantities to the nearest pint and unit prices to the nearest pence. Since most milk is sold in units of pints in England, 91% of the market share by volume in our sample, rounding has only a negligible effect on monthly expenditures. Finally, we extract a relatively homogenous subsample of couples with children, which gives us a working sample of 5,648 households.

Table 1 shows market shares by value and volume, and Table 2 displays unit prices (£/pint) for each type of milk.¹⁶ Semi-skimmed milk has more than 55% of the market share by both value and volume, followed by whole milk at about 35% and skimmed milk at about 8%. Whole and semi-skimmed milk are about £0.31/pint, and skimmed milk is slightly more expensive at about £0.32/pint, which means that there is little discernable price gradient with respect to fat content.

Note that implicit in our implementation is the assumption that milk characteristics are weakly separable from all other characteristics. This is perhaps contentious, but it is also necessary in order to make the implementation tractable. However, we note that the procedures developed in this paper can be extended to more complicated systems, which is of particular value in the nutritional context mentioned previously, where the effects of hypothetical own- and cross-price variation and any corresponding substitution patterns are of significant academic and policy interest. We leave this as a subject of future research.

Our application has $K = 3$ goods and $J = 2$ characteristics, and the technology \mathbf{A} is given by

$$\mathbf{A} = \begin{pmatrix} 1 & 20.869 \\ 1 & 9.429 \\ 1 & 0.568 \end{pmatrix}. \quad (11)$$

The three goods are whole, semi-skimmed, and skimmed milk. Following Blow, Browning, and Crawford (2008), the first characteristic is “milkiness”, which is the combination of water, calcium, vitamins, etc., that can be found in any milk. A pint of milk contains a pint of milkiness. The second characteristic is fat (grams/pint), which of course varies across types of milk. A pint of whole, semi-skimmed, and skimmed milk contains 20.869, 9.429, and 0.568 grams of fat on average by volume, respectively. Therefore, consuming a pint of each type of milk provides 30.866 grams of fat altogether. Note that since preferences over milkiness and fat are non-satiated, the structure allows fat to be a “bad” characteristic.

4 Results

The consistency results are shown in Table 3. Here we find that of 5,648 households, about 90.0% are q -rationalisable and 52.0% are z -rationalisable with imputed prices. As expected, the pass rate

¹³We exclude all organic, soya, UHT, speciality milks (e.g., buttermilk, chocolate milk, goats’ milk, etc.), and super fat milks. We also exclude promotional purchases, as well as cases with pack size discrepancies and expenditure anomalies. We then focus on the major brands, each of which has at least a 1% market share by volume in our sample.

¹⁴Aggregating by month effectively treats milk as a non-durable, non-storable good. This is unproblematic as long as milk is consumed relatively quickly. Note, as in Blow, Browning, and Crawford (2008), that aggregation is a simplification that potentially obscures container size and other relevant priced characteristics.

¹⁵Kantar partitions the UK into 10 geographical regions: London, Midlands, North East, Yorkshire, Lancashire, South, Scotland, Anglia, Wales and West, and South West. Imputing missing unit prices at a much finer regional level, i.e., postcode, is not always possible in the data.

¹⁶Note that these figures are based on observed data only.

for the goods model is higher than the pass rate for the characteristics model since z -rationalisation implies q -rationalisation. Similarly, about 97.7% of households are q -rationalisable and 88.3% are z -rationalisable with missing prices. Recall that the revealed preference conditions with missing prices are more permissive than with imputed prices. Our results are in line with the consistency results in Blow, Browning, and Crawford (2008).

The major contribution of this paper lies in identifying the model which best fits the data in light of these pass rates, and in order to do so, we consider the areas of the goods and characteristics models under different sets of priors. As shown in Table 4, $\bar{a}^q \approx 0.857$ and $\bar{a}^z \approx 0.000$ under (P1). Recall that dimension reduction often has the effect of reducing characteristics areas to zero under (P1). Consequently, $m(\bar{r}^q, \bar{a}^q) \approx 0.043$ and $m(\bar{r}^z, \bar{a}^z) \approx 0.520$ under (P1). For the goods model, this figure masks a great deal of heterogeneity. Among 5,648 households, 564 (10.0%) have negative predictive success, 2,750 (48.7%) have zero predictive success, and 2,334 (41.3%) have positive predictive success. The standard deviation of predictive success is 0.291. For the characteristics model, predictive success is completely determined by the pass rate. Lastly, $\delta(\Delta\bar{r}, \Delta\bar{a}) \approx 0.477$ under (P1). Once again, this figure obscures some important heterogeneity. Among 5,648 households, 1,203 (21.3%) have negative relative predictive success, 952 (16.9%) have zero relative predictive success, and 3,493 (61.8%) have positive relative predictive success. The standard deviation of relative predictive success is 0.538. The characteristics structure fits the data better on average, but for a sizable 38.2% of households the goods structure fits the data at least as well. Since these procedures are maximally heterogeneous, we later some explore logical associations with observable household attributes that may be useful in explaining some of this variation.

As previously noted, the tremendous power of the characteristics model relative to the goods model under (P1) is largely attributable to the dimension-reducing properties of the former. Under (P2), unobserved prices are set arbitrarily high, which restricts the number of feasible dimensions at the outset and therefore effectively controls for dimension reduction. As shown in Table 4, $\bar{a}^q \approx 0.948$ and $\bar{a}^z \approx 0.861$ under (P2), which implies that $\bar{a}^z - \bar{a}^q \approx -0.087$ is the average change in precision attributable to the characteristics structure alone. However, the average change in accuracy $\bar{r}^z - \bar{r}^q \approx -0.094$ under missing prices is slightly greater than the average change in precision, which means that $\delta(\Delta\bar{r}, \Delta\bar{a}) \approx -0.007$ under (P2). After controlling for dimension reduction, there is little between the two models, and even a slightly greater fit for the goods structure.¹⁷ This result suggests that in this instance much of the increase in power associated with the characteristics structure is attributable to its dimension-reducing capabilities rather than the particular technological mapping that we adopt.

As before, the (P2) results in Table 4 mask some important heterogeneity. For the goods and characteristics models, among 5,648 households, 131 and 104 (2.3% and 1.8%) have negative predictive success, 4,764 and 5,015 (84.3% and 88.8%) have zero predictive success, and 753 and 529 (13.3% and 9.4%) have positive predictive success, respectively. The standard deviations of predictive success are 0.175 and 0.153 in the goods and characteristics models, respectively. Furthermore, $\delta(\Delta\bar{r}, \Delta\bar{a}) \approx -0.007$ also obscures some heterogeneity. Among 5,648 households, 528 (9.3%) have negative relative predictive success, 4,799 (85.0%) have zero relative predictive success, and 321 (5.7%) have positive relative predictive success. The standard deviation of relative predictive success is 0.133. While the goods structure fits the data very slightly better on average, for a sizable 90.7% of households the characteristics fits the data at least as well.

Dimension reduction has such a pronounced effect on power in part because the outcome space is continuous. Under (P3), a discrete consumption space moderates these effects. As shown in Table 4, $\bar{a}^q \approx 0.743$ and $\bar{a}^z \approx 0.134$ under (P3). Consequently, $m(\bar{r}^q, \bar{a}^q) \approx 0.157$ and $m(\bar{r}^z, \bar{a}^z) \approx 0.386$ under (P3). For the goods and characteristics models, among 5,648 households, 440 and 181 (7.8% and 3.2%) have negative predictive success, 2,504 and 3,029 (44.3% and 53.6%) have zero predictive success, and 2,704 and 2,438 (47.9% and 43.2%) have positive predictive success. The standard deviations of predictive success are 0.322 and 0.476 in the goods and characteristics models, respectively. Lastly, $\delta(\Delta\bar{r}, \Delta\bar{a}) \approx 0.229$ under (P3). Among 5,648 households, 1,463 (25.9%) have negative relative predictive success, 1,332 (23.6%) have zero relative predictive success, and 2,853

¹⁷However, recall that for the characteristics model, it is sometimes possible to construct virtual prices that just allow for observed behaviour. This procedure also controls for dimension reduction, but just, and therefore necessarily increases the power of the characteristics model. As a result, 0.861 is an upper bound on \bar{a}^z , 0.022 is a lower bound on $m(\bar{r}^z, \bar{a}^z)$, and -0.007 is a lower bound on $\delta(\Delta\bar{r}, \Delta\bar{a})$ under (P2).

(50.5%) have positive relative predictive success. The standard deviation of relative predictive success is 0.529. The characteristics structure fits the data better on average, but for a sizable 49.5% of households the goods structure fits the data at least as well.

Altogether, the results suggest that the characteristics model fits the data better than does the goods model, but that dimension reduction rather than the technological mapping we adopt is the determining factor. The relative predictive success measure shows that the characteristics model performs relatively best under (P1), followed by (P3) and then (P2). The (P1) results very strongly support the characteristics structure, which is largely due to dimension reduction over a continuous outcome space. In order to mitigate against this *a priori* advantage, we control for dimension reduction under (P2) and find there to be very little between the two models. Lastly, we moderate the effects of dimension reduction by discretising the outcome space under (P3). One of our most important messages is that any claims about power and fit depend heavily upon feasibility priors, and also that it is important to isolate the changes in power and fit associated with dimension reduction from those associated with the particular structure of the model. In our example, had we observed positive relative predictive success under (P2), we might conclude that both dimension reduction *and* the particular characteristics structure we assume are important in explaining some of the observed variations in the data. However, in this instance, alternatives such as weak separability or other technological mappings from goods to characteristics may have performed better.

Here we briefly investigate via a descriptive regression analysis whether consistency, power, and fit are associated with observable household attributes. The controls include dummy variables corresponding to income bands and the main shopper being female, as well as continuous variables corresponding to the age of the main shopper in the household and the size of the household. Overall the statistical explanatory power of these control variables is incredibly weak, as suggested by low R^2 and pseudo- R^2 values, which leads us to conclude that consistency, power, and fit are largely unexplained by observable household attributes. However, we do observe a slight income gradient, whereby wealthier households are less likely to be consistent with either model, which disappears in the goods case but not the characteristics case after accounting for power. Furthermore, households with older main shoppers are more likely to be consistent with either model, but these effects are minimised after adjusting for power. The female main shopper dummy variable fails to give a significant result across many specifications. Lastly, the effect of household size has a positive and statistically significant effect on consistency and predictive success for the characteristics but not the goods model, and consequently, on relative predictive success for the characteristics model.

5 Conclusions

In this paper, we have compared the goods and characteristics models of the consumer within a nonparametric revealed preference framework. Of primary interest has been to make a comparison on the basis of predictive success that takes into account dimension reduction. This has allowed us to nonparametrically identify the model which best fit the data.

In the implementation on household panel data from the UK milk market, we have shown that on average the characteristics model fits the data better than does the goods model. Most importantly, we have shown that the better fit of the characteristics model is entirely attributable to dimension reduction. As a caution, we note that since the restrictive power of a model depends upon the feasible outcome space over which it is defined, any results must be interpreted in light of their corresponding feasibility priors. An analogy in statistical testing is that the result of the test depends upon the null hypothesis to be rejected. This should come as no surprise to econometricians, who are in the habit of constructing hypotheses in light of the possibility of Type I and Type II errors.

The methods developed in this paper can, in principle, be applied to other models of a similar class, to which we can exploit structure and power differentials and make a comparison on the basis of predictive success taking into account dimension reduction. Since the dimension-reducing property of the characteristics structure is the driving force behind its improved fit, there may be other dimension-reducing structures that provide even better fits. Applied microeconomists have long recognised the value of dimension reduction in empirical work. It simplifies the problem, both for the agent and the observer, and at the very least is implicitly invoked in nearly all applied

empirical work. Typically, separability restrictions or characteristics structures are introduced to reduce the number of dimensions, but it is *a priori* unclear how or where these restrictions should be imposed. Future research on dimension reduction should explore where it pays dividends in terms of fit to be strict versus lenient with structural assumptions. It should also explore the effects of hypothetical price variation in order to bound comparative statics of interest.

In the implementation, we have highlighted a relevant nutritional application. The procedures developed in this paper have allowed us to nonparametrically identify the relative importance of nutrients in explaining food consumption. Economists are often interested in estimating theoretically consistent demand systems for food. Since food contains some nutrients that are healthy in small amounts and unhealthy in large amounts, the characteristics model is a very plausible structure for food and nutrients. Individuals may trade off between taste and health, which is of particular interest for goods that we may perceive to ultimately become “bads”. The procedures developed in this paper can be extended to more complicated demand systems, which is of particular value in this nutritional context, where the effects of hypothetical own- and cross-price variation and corresponding substitution patterns are of significant academic and policy interest. If consumers have preferences over nutrients rather than food, then the implications of price variation change dramatically. We leave this as a subject of future research. There are of course other markets to explore, e.g., housing, automobile, education, and labour, which may be of great interest to applied researchers and policymakers, which again we leave for future work.

This paper has neglected any considerations on the supply side of the economy. However, applied microeconometricians working on problems of industrial organisation are very often interested in both sides of the market. In the context of the characteristics framework, prices are endogenous, and hence there is an equilibrium model underlying hedonic prices. The basic equilibrium theory posits that individual consumers and differentiated producers make choices in characteristics space, the results of which are market-clearing equilibrium prices that reflect implicit intrinsic mappings from the prices of goods to the characteristics they embody. In order to apply the results of this paper to applied microeconomic work in these areas, we must further explore the driving forces behind identification in equilibrium hedonic models under minimal assumptions.

References

- [1] Afriat, S.N. (1967). ‘The construction of utility functions from expenditure data’, *International Economic Review*, vol. 8(1), pp. 67–77.
- [2] Andreoni, J. and Harbaugh, W.T. (2008). ‘Power indices for revealed preference tests’, mimeo, University of California, San Diego.
- [3] Banks, J., Blundell, R. and Lewbel, A. (1997). ‘Quadratic Engel curves and consumer demand’, *Review of Economics and Statistics*, vol. 79(4), pp. 527–539.
- [4] Beatty, T.K.M. and Crawford, I.A. (2011). ‘How demanding is the revealed preference approach to demand?’, *American Economic Review*, vol. 101(6), pp. 2782–2795.
- [5] Berry, S., Levinsohn, J. and Pakes, A. (1995). ‘Automobile prices in market equilibrium’, *Econometrica*, vol. 63(4), pp. 841–890.
- [6] Blow, L., Browning, M. and Crawford, I. (2008). ‘Revealed preference analysis of characteristics models’, *Review of Economic Studies*, vol. 75(2), pp. 371–389.
- [7] Bronars, S.G. (1987). ‘The power of nonparametric tests of preference maximization’, *Econometrica*, vol. 55(3), pp. 693–698.
- [8] Deaton, A. and Muellbauer, J. (1980). ‘An almost ideal demand system’, *American Economic Review*, vol. 70(3), pp. 312–326.
- [9] Diewert, W.E. (1973). ‘Afriat and revealed preference theory’, *Review of Economic Studies*, vol. 40(3), pp. 419–425.

- [10] Fostel, A., Scarf, H.E. and Todd, M.J. (2004). ‘Two new proofs of Afriat’s theorem’, *Economic Theory*, vol. 24(1), pp. 211–219.
- [11] Gorman, W.M. (1956). ‘A possible procedure for analysing quality differentials in the egg market’, *Review of Economic Studies*, vol. 47(5), pp. 843–856.
- [12] Griffith, R. and Nesheim, L. (2010). ‘Estimating households’ willingness to pay’, mimeo, University College London.
- [13] Griffith, R. and O’Connell, M. (2009). ‘The use of scanner data for research into nutrition’, *Fiscal Studies*, vol. 30(3–4), pp. 339–365.
- [14] Houthakker, H.S. (1950). ‘Revealed preference and the utility function’, *Economica*, vol. 17(66), pp. 159–174.
- [15] Lancaster, K.J. (1966). ‘A new approach to consumer theory’, *Journal of Political Economy*, vol. 74(2), pp. 132–157.
- [16] Leicester, A. and Oldfield, J. (2009). ‘Using scanner technology to collect expenditure data’, *Fiscal Studies*, vol. 30(3–4), pp. 309–337.
- [17] McFadden, D. (1973). ‘Conditional logit analysis of qualitative choice behavior’, in (P. Zarembka, ed.), *Frontiers in Econometrics*, pp. 105–142, New York: Academic Press.
- [18] Polisson, M. and Quah, J. (2011). ‘Discreteness, separability, and revealed preference’, mimeo, University of Oxford.
- [19] Richter, M.K. (1966). ‘Revealed preference theory’, *Economica*, vol. 34(3), pp. 635–645.
- [20] Samuelson, P.A. (1938). ‘A note on the pure theory of consumer’s behaviour’, *Economica*, vol. 5(17), pp. 61–71.
- [21] Samuelson, P.A. (1948). ‘Consumption theory in terms of revealed preference’, *Economica*, vol. 15(60), pp. 243–253.
- [22] Selten, R. (1991). ‘Properties of a measure of predictive success’, *Mathematical Social Sciences*, vol. 21(2), pp. 153–167.
- [23] Varian, H.R. (1982). ‘The nonparametric approach to demand analysis’, *Econometrica*, vol. 50(4), pp. 945–973.
- [24] Varian, H.R. (1983). ‘Non-parametric tests of consumer behaviour’, *Review of Economic Studies*, vol. 50(1), pp. 99–110.

Table 1: Market shares

Market share (%)	Whole	Semi-skimmed	Skimmed
By volume	35.2	56.8	8.0
By value	35.0	56.7	8.3

Table 2: Prices

Price (£/pint)	Whole	Semi-skimmed	Skimmed
Mean	0.31	0.31	0.32
Median	0.29	0.29	0.32
Std	0.04	0.04	0.04

Table 3: Consistency results

	Imputed prices		Missing prices	
	<i>q</i> -rationalisation	<i>z</i> -rationalisation	<i>q</i> -rationalisation	<i>z</i> -rationalisation
Pass	5,084 (90.0%)	2,935 (52.0%)	5,516 (97.7%)	4,987 (88.3%)
Fail	564 (10.0%)	2,713 (48.0%)	132 (2.3%)	661 (11.7%)

Table 4: Power results

Priors	\bar{r}^q	\bar{r}^z	\bar{a}^q	\bar{a}^z	$m(\bar{r}^q, \bar{a}^q)$	$m(\bar{r}^z, \bar{a}^z)$	$\delta(\Delta\bar{r}, \Delta\bar{a})$
P1	0.900	0.520	0.857	0.000	0.043	0.520	0.477
P2	0.977	0.883	0.948	0.861	0.029	0.022	-0.007
P3	0.900	0.520	0.743	0.134	0.157	0.386	0.229

Figure 1a: Restricted characteristics budget set




Figure 1b: Unrestricted characteristics budget set




Figure 2: Example of a q -rationalisation




Figure 3: Example of an area




Figure 4: Example of a z -rationalization

