Revealed Preference and Welfare Analysis

Eileen Liong Tipoe

DPhil Economics

Christ Church College

Acknowledgements: The author would like to thank Ian Crawford for his guidance and supervision, and Abigail Adams for helpful comments and suggestions.

Revealed Preference and Welfare Analysis
Eileen Liong Tipoe
DPhil Economics
Christ Church College
Trinity Term 2017

Abstract

This thesis uses nonparametric revealed preference methods to derive new tests for consistency with models of consumer behaviour, and discuss the implications for welfare analysis. Chapter 1 demonstrates how to conduct revealed preference analysis when prices, and hence budget constraints, are only partially observed. This chapter extends the revealed preference results of Crawford and Polisson (2015), derived for the static case, to dynamic settings, allowing for storability of goods. Necessary and sufficient conditions for consistency with intertemporal models are derived, which do not require the researcher to distinguish between corner solutions and unavailability of the good, or to impute prices. Chapter 2 discusses the validity of using reported happiness measures as proxies of utility or social welfare, by testing for consistency between revealed and reported preference orderings in Japanese household survey data. Although the expenditure behaviour of most households is consistent with standard models of utility maximisation, it is generally inconsistent with the preference ordering given by their reported happiness. This inconsistency is likely due to reporting error in the happiness measure, and suggests that happiness and utility are empirically distinct and non-interchangeable. Chapter 3 investigates the effect of price inattention on inflation misperceptions and cost-of-living indices, by developing a behavioural model in which consumers only notice price changes above a certain threshold. A data application, using supermarket scanner data, demonstrates that this model generates plausible results; in particular, consumers have more accurate perceptions of inflation during periods of high or volatile inflation, but may substantially misperceive inflation when it is low. These results have important implications for conducting welfare analysis when consumers are not fully attentive to price changes.

Contents

Introduction

This thesis focuses on methods and challenges in recovering preferences from behaviour. Revealed preference methods provide a link between observable behaviour and the unobservable preferences that generate it, allowing researchers to determine whether behaviour is consistent with a particular model and recover the set of possible underlying preferences. One challenge for preference recovery and welfare analysis is when the researcher cannot observe the choices that an individual could have made, which is the case when some price data are missing. Even when preference recovery is possible, researchers can only infer the preference ordering and not the degree to which one choice is preferred over another. However, it may be possible to put additional structure on preferences by using cardinal information on well-being such as happiness reports. Behavioural models also pose additional challenges for preference recovery, because unlike in neoclassical models, preferences are not always time-consistent or stable. These three issues form the main discussion of this thesis: Chapter 1 shows that preference recovery is still possible when budget constraints are not always observed, Chapter 2 discusses the usefulness of happiness measures for preference recovery, and Chapter 3 shows how to recover preferences and conduct welfare analysis when consumers are not fully attentive to price changes.

The revealed preference approach, developed by Afriat (1967) and Varian (1982), uses observed behaviour to construct unobservable elements of economic decision making, such as indifference curves and preferences. Another object of interest is testing whether behaviour is consistent with specific economic models, using observable data only. This approach is nonparametric, so methods used do not require any untestable assumptions such as the statistical distributions of unobserved variables or the relationships between them. Revealed preference tests therefore check for consistency with the model's assumptions alone, rather than jointly testing the validity of the model and its auxiliary assumptions. These tests usually involve systems of inequalities that can be verified using linear or integer programming methods.

Consistency usually requires the data to satisfy particular conditions, which are observable consequences of that model. If these conditions are also sufficient, then it is possible to construct unobservable variables and parameters of the given model from observable data[1]. Given a dataset with prices and quantities, the Generalised Axiom of Revealed Preference (GARP) is one way to check whether this data is rationalisable, meaning that it could have been generated by maximising a well-behaved utility function:

Definition 0.1.[2] The data $\mathbf{D} = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,\dots,T}$ **satisfies GARP** if and only if we can construct preference relations R_0 and R such that for all $r, s, t, u, v \in T$:

(1) If $\mathbf{p}_t\mathbf{q}_t \geq \mathbf{p}_t\mathbf{q}_s$, then $\mathbf{q}_t R_0 \mathbf{q}_s$.

(2) If $\mathbf{q}_r R_0 \mathbf{q}_s$, $\mathbf{q}_s R_0 \mathbf{q}_t$,... and $\mathbf{q}_u R_0 \mathbf{q}_v$, then $\mathbf{q}_r R \mathbf{q}_v$.

(3) If $\mathbf{q}_t R \mathbf{q}_s$, then $\mathbf{p}_s\mathbf{q}_s < \mathbf{p}_s\mathbf{q}_t$.

The preference relations R_0 and R stand for 'directly revealed preferred' and 'revealed preferred' respectively. (1) is a condition on R_0: for observed choices to be optimal, consumption bundle $\mathbf{q}_t$ must be directly revealed preferred to $\mathbf{q}_s$ if $\mathbf{q}_t$ was chosen when $\mathbf{q}_s$ was also affordable (otherwise the consumer could do better by choosing $\mathbf{q}_s$, given his budget constraints). Likewise, (3) is a condition on R: if consumption bundle $\mathbf{q}_t$ is revealed preferred to $\mathbf{q}_s$, then $\mathbf{q}_t$ must cost at least as much as $\mathbf{q}_s$. Condition (2) states that R must obey transitivity: if consumption bundle A is preferred to B, B preferred to C,..., and Y pre-

[1]In parametric approaches, it is possible to uniquely identify these unobservable variables. Conversely, with the non-parametric approach, there may be many preference orderings or a range of model parameters that are consistent with the observed behaviour. However, the focus of revealed preference tests is to determine whether a given model could generate the observed behaviour, and not to determine the actual values of unobservables.

[2]Paraphrased from Varian (1982).

ferred to Z, then A must be preferred to Z in order to avoid an illogical cycle of preferences. Constructing preference relations over data is computationally feasible but may not be practical for a very large number of observations. However, there are computationally easier ways to check for consistency with utility maximisation, which are stated in Afriat's Theorem:

Theorem 0.1 (Afriat's Theorem)[3]. Given the dataset $\mathbf{D} = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$, the following conditions are equivalent:

(1) The dataset $\mathbf{D}$ is consistent with utility maximisation, and therefore there exists a well behaved utility function[4] that could have generated the data.

(2) The dataset $\mathbf{D}$ satisfies the Generalised Axiom of Revealed Preference (GARP).

(3) There exists numbers $u_t \in \mathbb{R}$ and $\lambda > 0$ such that

$$u_s \leq u_t + \lambda_t \mathbf{p}_t \cdot (\mathbf{q}_s - \mathbf{q}_t) \text{ for all } s, t \in T.$$

(4) There exist numbers $y_{st} \in 0, 1$ such that

(i) $\mathbf{p}_t(\mathbf{q}_t - \mathbf{q}_s) \leq \mathbf{p}_t \cdot \mathbf{q}_t(1 - y_{st})$

(ii) $y_{su} + y_{ut} \leq 1 + y_{st}$

(iii) $\mathbf{p}_s(\mathbf{q}_s - \mathbf{q}_t) < \mathbf{p}_s \cdot \mathbf{q}_s y_{st}$

for all $u, s, t \in \{1, ..., T\}$, where $y_{st} = 1$ if $\mathbf{q}_s R \mathbf{q}_t$.

Condition (4), the integer-program equivalent of GARP, gives a set of inequalities that are linear when conditioning on $\{y_{st}\}_{s,t=1,...,T}$ (which define the preference relations between time periods), and can therefore be solved using linear programming methods[5]. Afriat's Theorem states that finding a solution to either (3) or (4) is sufficient to conclude that the

[3] Paraphrased from Varian (1982) and Cherchye et al (2011a); see these papers for proof.

[4] Assuming that the utility function is well-behaved (nonsatiated, concave, continuous, and monotonic) is without loss of generality, because with finite data and linear prices it is impossible to detect non-concavities (Afriat, 1967).

[5] Since the preference relations are unknown, all possible combinations of $\{y_{st}\}_{s,t=1,...,T} \in \{0, 1\}$ must be checked. Also note that the inequalities may be consistent with more than one set of preference relations, so if a solution exists then it may not be unique.

data is consistent with static utility maximisation. Throughout this thesis, the linear programming test (3) and integer programming tests of the form (4) will be both be referred to as 'GARP tests'.

These revealed preference conditions can be extended to intertemporal models, although the assumption of perfect foresight is necessary because of difficulties dealing with uncertainty. Browning (1989) first presents revealed preference conditions for the life cycle model, and Crawford (2010) and Demuynck and Verriest (2013) extend these conditions to allow for habit formation and other forms of consumption dependence. Chapter 1 derives revealed preference conditions for these models when prices are only partially observed.

Although Afriat's Theorem applies to individual-level data, there are also revealed preference conditions for models of household behaviour, which are applicable to household-level data such as expenditure surveys. Cherchye et al (2008) and Cherchye et al (2011b) derive tests for these models, which are discussed in Chapter 2.

All of these revealed preference tests give a binary result, meaning that data either 'Pass' (behaviour is consistent with utility maximisation according the given model) or 'Fail' (behaviour is not consistent with the model). However, inconsistencies could arise because of measurement error in survey data or optimisation error. The Afriat efficiency index (Afriat, 1972) measures how 'close' the data are to passing a test and allows researchers to determine whether optimisation error is a plausible reason for the inconsistency.

Revealed preference tests are conducted on each individual or household separately, and so allow for any type of heterogeneity in preferences across the sample. The simplest way to evaluate a model's ability to explain observed behaviour is the pass rate, which is the proportion of individuals or households in a given dataset that pass the associated revealed preference test. However, a high pass rate alone does not necessarily mean the model is 'good'; instead it could indicate that the model does not put many observable restrictions on behaviour. There are many ways of measuring the restrictiveness of models, the most popular method being the power calculation developed by Bronars (1987), which indicates the ability of the revealed preference test to detect violations of the model. The measure of predictive success developed by Beatty and Crawford (2011) evaluates how 'good' a model is at explaining behaviour by comparing the pass rate with power. These three measures (pass rate, power, predictive success) will be used throughout this thesis to interpret revealed preference test results.

1. Revealed Preference Conditions with Partially Observed Prices

1.1 Motivation

Datasets of consumer purchases can contain missing observations; in particular, when a consumer does not purchase a good, there is no information about the price at which the consumer could have bought it. So for a given consumer that can purchase K goods $(q^1, ..., q^K)$ over T time periods, instead of observing the dataset $\mathbf{D}^* = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$, the researcher only observes the dataset $\mathbf{D} = \{p_t^k | q_t^k > 0, q_t^k\}_{t=1,...,T}^{k=1,...,K}$. With fully observed prices, the researcher can construct the consumer's budget constraints and make some statements about the consumer's preferences over bundles of goods. Partially observed prices is a problem for revealed preference analysis, since the researcher cannot construct the budget set of a particular consumer and hence cannot determine the consumer's feasible alternatives. One particular setting in which this issue is prevalent is highly disaggregated scanner data, which is the focus of this chapter.

Zero demand for a good can be observed for many reasons, including rationing (the good may be unavailable at time of purchase), the consumer's choice (corner solution), and infrequent purchase (over a short time period, no purchase may be recorded for goods that are bought infrequently). For household expenditure surveys, zero purchases may also be due to misreporting by the consumer: if the survey was given long after the goods were purchased, the consumer may not remember the exact time of purchase and so may report a zero even if the good was actually purchased in the given time frame.

Existing literature has come up with ways to deal with missing prices, depending on the reason. If all missing prices are due to corner solutions, then the standard Tobit model, which accounts for the left-censoring in the data, can be used. For household expenditure surveys, which may suffer from misreporting, Deaton and Irish (1984) suggest a test to determine whether or not reporting error is present. Meghir and Robin (1992) and Keen (1986) extend the standard Tobit model to address the issue of infrequent purchasing in order to consistently estimate both linear and nonlinear Engel curves, though the actual number of purchases needs to be observed for identification. The fundamental problem with these methods is that the way to deal with missing prices depends on the reason, which is typically unknown to the researcher without additional information such as the availability of goods at the time of purchase. These methods also rely on independence or distributional assumptions such as normality, and assumptions on the functional form of the utility function, not all of which are testable. Any revealed preference test using these methods therefore becomes a joint test of maximising behaviour and the assumptions required for consistent estimation, rather than a test of maximising behaviour alone.

Missing prices can be imputed, for example by taking an average of observed prices for the same good purchased by similar consumers in the same location and day (or similar goods/

locations/days), but this method can be difficult to implement if the range of goods, locations, and days is very large. With a highly disaggregated panel such as supermarket scanner data, the curse of dimensionality can make this method infeasible in practice. In settings where goods are sometimes unavailable, the imputed price also depends on the reason why the good was not purchased: it should either be the reservation price if the consumer did not consider it optimal to purchase it, or the virtual price if the good was out of stock. However, the researcher usually does not observe the reason for a zero purchase, so any imputation will be somewhat ad-hoc. Another problem with imputing prices is that rationalisability and welfare analysis may be sensitive to the method chosen. In particular, for a given dataset with partially observed prices, it is possible for the dataset to pass GARP for one set of imputed prices, but fail GARP for a different set of imputed prices[6]. Therefore, it would be desirable to test for rationality without needing to impute prices.

Revealed preference approach

The revealed preference approach, developed by Afriat (1967) and Varian (1982), is nonparametric and allows researchers to test whether or not a dataset is consistent with utility-maximising behaviour according to a particular model without having to impose additional structure on the utility function. Crawford and Polisson (2015) derive revealed preference conditions for the static case where missing prices are due to either rationing or a corner solution[7]:

Proposition 1.1.[8] The dataset $\mathbf{D} = \{p_t^k | q_t^k > 0, q_t^k\}_{t=1,...,T}^{k=1,...,K}$ is rationalisable by a well-behaved (nonsatiated, concave, continuous, and monotonic) utility function $u : \mathbb{R}_+^K \to \mathbb{R} \Leftrightarrow$ For $\forall t = 1, ..., T$, there exist numbers $u_t \in \mathbb{R}$ and $\lambda_t \in \mathbb{R}_{++}$, and vectors $\hat{\boldsymbol{\pi}}_t \in \mathbb{R}_{++}^K$ such that

[6] See Crawford and Polisson (2015) for some examples.

[7] Aside from these two reasons and infrequent purchase (storability), this chapter will not deal with other reasons for missing prices, such as misreporting.

[8] Adapted from Proposition 1, Crawford and Polisson (2015).

$$u_s \leq u_t + \lambda_t \hat{\boldsymbol{\pi}}_t'(\mathbf{q}_s - \mathbf{q}_t) \text{ for } \forall s, t = 1, ..., T$$

$$\hat{\pi}_t^k = p_t^k \text{ for any } q_t^k > 0 \text{ for } \forall k = 1, ...K, t = 1, ..., T$$

where λ_t is the the marginal utility of income at observation t and p_t^k is the non-discounted (actual) price of good k at observation t.

According to Proposition 1.1, the existence of a solution to this system of T distinct inequalities (numbers $\{u_t, \lambda_t, \hat{\pi}_t^k\}_{t=1,...,T}^{k=1,...,K}$) is the only requirement for the dataset $\mathbf{D}$ to be consistent with static utility maximisation. This result is a generalisation of the inequalities in Afriat's Theorem to allow $q_t^k = 0$ for some k, t. The vector of known actual prices is replaced by a vector of support prices, which are unknowns for goods not purchased but (as stated in the additional condition) are equal to the actual price for goods that are purchased. If all prices are observed, then Proposition 1.1 reduces to Afriat's Theorem. Instead of imputing missing prices, this revealed preference test involves finding a set of numbers (support prices) that, when used in place of the missing prices, rationalises the data[9]. Therefore in order to conduct the revealed preference test, there is no need for the researcher to impute missing prices or know whether missing prices are due to rationing or a corner solution.

This chapter will extend the results of Crawford and Polisson (2015) to allow for intertemporal choice (utility maximisation, given lifetime income and discounted prices) and storability of goods (having an unobserved stock of goods or receiving utility from a good in the periods after purchase), both of which would affect purchasing decisions in any given time period. Revealed preference conditions for dynamic models with these characteristics have been derived by Browning (1989), Crawford (2010), and Demuynck and Verriest (2013), but only for complete datasets. The results of this chapter are of practical importance, because most consumer data is in panel format and consumers may make consumption plans over multiple periods rather than per-period purchasing decisions.

[9] Remember that $\hat{\boldsymbol{\pi}}_t$ need not be unique; there may be a set of support prices that rationalise the data. The revealed preference test checks whether or not this set is empty.

The remainder of this chapter is organised as follows: Section 1.2 incorporates intertemporal choice and derives revealed preference conditions for the Browning (1989) life cycle model, Section 1.3 incorporates storability via the habits models of Crawford (2010) and Demuynck and Verriest (2013) and shows that the general conclusions do not depend on the particular way that storability is modeled, Section 1.4 demonstrates that these revealed preference conditions are falsifiable by presenting simple examples of datasets that fail each of the models from Sections 1.2 and 1.3, Section 1.5 concludes.

1.2 Adding dynamics: The life cycle model

As discussed in Section 1.1, adding dynamics makes the model of consumer behaviour more realistic, as current decisions often need to take future variables such as next period's income or prices into account. The Browning (1989) life cycle model used in this section allows consumers to base their purchasing decisions on both current and future prices of goods, by having a lifetime budget constraint instead of a per-period budget constraint. This life cycle model assumes that consumers (1) have perfect foresight[10], (2) face perfect capital markets, and (3) have additively separable and stable preferences across time periods (consumption independence). Each consumer also has a discount factor, $\beta \in (0, 1]$, which represents the value they assign to future consumption relative to this period's consumption[11]. It can either be assumed that the consumer is infinitely lived (as in Demuynck and Verriest, 2013) or lives for $T < \infty$ periods (as in Browning, 1989); both assumptions result in the same revealed preference conditions, so this chapter takes $T = \infty$ without loss of generality.

[10]This chapter does not incorporate uncertainty, since revealed preference conditions have not been formally derived under uncertainty even for complete datasets.

[11]Although it is theoretically possible for $\beta > 1$, there is very little empirical evidence that consumers value the future more than the present. In a review of existing studies that estimate discount factors, Frederick, Loewenstein, and O'Donoghue (2002) find that $\beta \in (0, 1]$ in all studies regardless of the time horizon considered (ranging from 1 week or less to 15 years). Also, while the revealed preference conditions derived are still testable under the assumption that $\beta \in (0, \infty)$, it is computationally convenient to assume that β has an upper limit. For these two reasons, this chapter will take $\beta \in (0, 1]$.

Each consumer's preferences over time periods can therefore be represented by the utility function

$$\sum_{t=1}^{\infty} \beta^{t-1} u(\mathbf{q}_t)$$

where $\mathbf{q}_t \in \mathbb{R}_+^K$ is a vector of K goods purchased in period t and the instantaneous utility function $u(\mathbf{q}_t)$ is assumed to be concave, continuous, and strictly increasing.

One way to identify the life cycle model is to assume the Strong Rational Expectations Hypothesis (SREH), which states that consumers aim to keep the marginal utility of discounted expenditure equal across periods i.e. $\lambda_t = \lambda$ for $\forall t$ (Bewley, 1977; Hall 1978). The SREH follows from assumptions (1)-(3) discussed above, and from having one lifetime budget constraint. Under the SREH, the consumer optimally allocates expenditure across time periods: if the marginal utility of discounted expenditure is lower in period t than period $t-1$, then the consumer can be made better off by transferring some expenditure from period t to period $t-1$. The SREH is plausible over short time periods or in the absence of uncertainty since consumers can properly account for price and income variability, though it is quite a strong assumption for long periods or if consumers face uncertainty. However, weaker versions of the SREH exist, for example allowing perfect foresight for only one or $T' < T$ time periods ahead (see Browning, 1989). The revealed preference conditions under the SREH generalise under weaker assumptions and revealed preference tests can be conducted in a similar way.

The consumer's maximisation problem is:

$$\max_{\mathbf{q}_t \in \mathbb{R}_+^K} \sum_{t=1}^{\infty} \beta^{t-1} u(\mathbf{q}_t) \text{ s.t. } \sum_{t=1}^{\infty} \boldsymbol{\rho}_t' \mathbf{q}_t = Y \text{ and } q_t^k = 0 \text{ for any } k \in \kappa_t$$

where $\boldsymbol{\rho}_t$ is the $K \times 1$ vector of discounted prices ($\rho_t^k = \frac{p_t^k}{\prod_{s=2}^{s=t}(1+i_s)}$), i_s is the interest rate in period s, p_t^k is the actual price of good k in period $t \in \{1, ...T\}$, Y is discounted total lifetime income (taken to be exogenous), and κ_t is the set of goods which are unavailable to the consumer in period t. Under the perfect foresight asssumption, the consumer knows

$\{p_t^k, i_t\}_{t=1,\dots,T}^{k=1,\dots,K}$ and Y at $t = 1$. The key difference between the static model from Section 1.1 and this dynamic model is the lifetime budget constraint instead of the per-period budget constraint, and the introduction of discounting (via the discount factor β and the use of discounted prices $\boldsymbol{\rho}_t$ instead of actual prices $\mathbf{p}_t$).

The lagrangian is therefore: $\mathcal{L} = \sum_{t=1}^{\infty} \beta^{t-1} u(\mathbf{q}_t) + \lambda(Y - \sum_{t=1}^{\infty} \boldsymbol{\rho}_t' \mathbf{q}_t) + \gamma_t^k(0 - q_t^k)$ and the first-order conditions (FOCs) are[12]:

$$\beta^{t-1} \frac{\partial u(\boldsymbol{q}_t)}{\partial q_t^k} = \lambda \rho_t^k \text{ if } q_t^k > 0$$

$$\beta^{t-1} \frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} \leq \lambda \rho_t^k \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t$$

$$\beta^{t-1} \frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \lambda \rho_t^k + \gamma_t^k \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t$$

where $\gamma_t^k > 0$ is the marginal cost to utility due to rationing good k in period t (the disutility from not being able to purchase the good). Note that in a given time period, not all goods may be purchased so there may be missing price observations. Since econometricians cannot observe either Y or κ_t given the dataset $\mathbf{D} = \{p_t^k | q_t^k > 0, q_t^k\}_{t=1,\dots,T}^{k=1,\dots,K}$, without additional data on the consumer it is impossible to determine whether demand for a good is zero due to the consumer's choice (a corner solution) or unavailability of the good. However, the following discussion shows that it is still possible to establish necessary and sufficient conditions for optimising behaviour when some price observations are missing, without having to know or make assumptions about the reason for zero demand.

This rationed maximisation problem can be expressed as an unrationed maximisation problem by using support prices instead of discounted prices. **Discounted support prices** ($\{\pi_t^k\}_{t=1,\dots,T}^{k=1,\dots,K}$) are strictly positive real numbers that generate exactly the same demand choices when used in the unrationed maximisation problem as those made in the rationed maximisation problem.

[12]In real-life situations, goods usually come in discrete amounts so $u(\mathbf{q}_t)$ may not be differentiable with respect to all goods. However, Polisson and Quah (2013) show that Afriat's Theorem still applies in a discrete consumption space, as long as the consumer's utility is increasing in a continuous good, which can be outside the subset of goods observed. Therefore, assuming that $u(\mathbf{q}_t)$ is differentiable is without loss of generality.

The support prices for this maximisation problem are:

$$\pi_t^k = \begin{cases} \rho_t^k & q_t^k > 0 \\ \rho_t^k + \frac{\gamma_t^k}{\lambda} & q_t^k = 0; k \in \kappa_t \\ \frac{\partial u(q_t)}{\partial q_t^k}\frac{1}{\lambda} & q_t^k = 0; k \notin \kappa_t \end{cases}$$

For purchased goods, support prices are actual discounted prices; for unavailable goods, support prices are the lowest prices that result in zero demand without the rationing constraint (virtual prices); for goods available but not purchased, support prices are the prices that make the consumer exactly indifferent between purchasing and not purchasing (reservation prices). The assumptions made earlier on the utility function (concavity, continuity, and monotonicity) are sufficient for the existence of support prices $\pi_t^k > 0$ for $\forall k, t$ that are consistent with any set of observed rationed demands[13]. The results presented in Neary and Roberts (1980) mean that any observed choices made when facing observed prices and rationing are exactly those made when facing support prices without rationing; thus rationed demand changes can be expressed in terms of unrationed demand changes. Using these support prices instead of observed prices, the solution to the rationed maximisation problem above is equivalent to that of the following unrationed maximisation problem:

$$\max_{\mathbf{q}_t \in \mathbb{R}_+^K} \sum_{t=1}^{\infty} \beta^{t-1} u(\mathbf{q}_t) \text{ s.t. } \sum_{t=1}^{\infty} \boldsymbol{\pi}_t^{'} \mathbf{q}_t = Y.$$

An important issue is whether there are necessary and sufficient conditions based only on observables for the data to be consistent with the life cycle model, even when some prices are missing. Definition 1.1 states what it means for the data to be consistent with the life cycle model:

Definition 1.1.[14] A utility function $u(.)$ $(u : \mathbb{R}_+^K \rightarrow \mathbb{R})$ **i-rationalises** the dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,\dots,T}$ if there exist support prices $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^K$ (where $\pi_t^k = \rho_t^k$ if $q_t^k > 0$) and

[13] This result is proved formally in Neary and Roberts (1980).
[14] Adapted from Crawford & Polisson (2015)

$\beta \in (0,1]$ such that for $\forall\ t \in \{1, ..., T\}$, $u(\mathbf{q}_t) \geq u(\mathbf{q})$ for any $\mathbf{q} \in \{\mathbf{q} \in \mathbb{R}^k_+ : \boldsymbol{\pi}_t'\mathbf{q} \leq \boldsymbol{\pi}_t'\mathbf{q}_t\}$.

This definition means that the observed choices are consistent with utility maximisation if we can find a utility function and support prices (actual prices for goods purchased, and reservation or virtual prices for goods not purchased) such that the observed behaviour is optimal i.e. is the solution to the unrationed maximisation problem. Due to concavity, continuity, and monotonicity of the utility function, the first-order conditions are sufficient for behaviour to be utility-maximising. Therefore, for the data to be consistent with a given model, there must be a utility function whose derivatives satisfy the first-order conditions of that model.

Definition 1.1 applied in this context implies that the dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ (where $q_t^k = 0$ for some $t \in \{1, ..., T\}$, $k \in \{1, ..., K\}$) is rationalisable by the life cycle model (i-rationalisable) if there exists a locally non-satiated, continuous, concave, and differentiable utility function $u(\mathbf{q}_t)$ and numbers $\lambda > 0$, $\gamma_t^k > 0$, $\beta \in (0,1]$ such that

$$\beta^{t-1}\frac{\partial u(\boldsymbol{q}_t)}{\partial q_t^k} = \lambda \rho_t^k \text{ if } q_t^k > 0$$

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} \leq \lambda \rho_t^k \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t$$

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \lambda \rho_t^k + \gamma_t^k \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t$$

While it seems as if the reason for missing prices matters since the first-order conditions to satisfy depend on the availability of the good, the proof of Theorem 1.1 shows that the data can be rationalised by treating all goods as if they were rationed or as if they were all corner solutions, or a mixture of the two. In this sense, the reason for missing prices does not matter; in both cases the first-order conditions are satisfied if we can find numbers (representing the marginal utility of discounted total lifetime expenditure, the discount factor, and the marginal disutility of rationing) that can rationalise the observed zero demands. Theorem 1.1 formalises this result.

Theorem 1.1 (Life cycle model). The following statements are equivalent:

(1) The dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ is i-rationalisable.

(2) There exist numbers $u_t \in \mathbb{R}$, $\lambda > 0$, $\beta \in (0, 1]$, and vectors $\boldsymbol{\pi}_t \in \mathbb{R}^K_{++}$ such that

$$u_s \leq u_t + \tfrac{1}{\beta^{t-1}}[\lambda\boldsymbol{\pi}_t'(\mathbf{q}_s - \mathbf{q}_t)] \text{ for } \forall s, t = 1, ..., T$$

$$\pi_t^k = \rho_t^k \text{ for any } q_t^k > 0 \text{ for } \forall k = 1, ...K, t = 1, ..., T.$$

***Proof*:**

(1) $\implies$ (2) [*Necessity*] The dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ is rationalisable by the life cycle model if there exists a locally non-satiated, continuous, concave, and differentiable utility function $u(.)$ and numbers $\lambda > 0$, $\gamma_t^k > 0$, $\beta \in (0, 1]$ such that

$$\tfrac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \tfrac{1}{\beta^{t-1}}\lambda\rho_t^k \text{ if } q_t^k > 0 \text{ ; } t = 1, ..., T, k = 1, ..., K$$

$$\tfrac{\partial u(\mathbf{q}_t)}{\partial q_t^k} \leq \tfrac{1}{\beta^{t-1}}\lambda\rho_t^k \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t \text{ ; } t = 1, ..., T, k = 1, ..., K$$

$$\tfrac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \tfrac{1}{\beta^{t-1}}[\lambda\rho_t^k + \gamma_t^k] \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t\text{; } t = 1, ..., T, k = 1, ..., K$$

By concavity of $u(.)$,

$$u_s \leq u_t + \nabla u(\mathbf{q}_t)'(\mathbf{q}_s - \mathbf{q}_t) \text{ for } \forall s, t = 1, ..., T$$

For $t = 1, ..., T, k = 1, ..., K$, define $\pi_t^k = \rho_t^k$ if $q_t^k > 0$ or $q_t^k = 0$ and $k \notin \kappa_t$, and $\pi_t^k = \rho_t^k + \frac{\gamma_t^k}{\lambda}$ if $q_t^k = 0$ and $k \in \kappa_t$. Note that $\pi_t^k > 0$ since ρ_t^k, γ_t^k, and $\lambda > 0$, so $\nabla u(\mathbf{q}_t) \leq \frac{1}{\beta^{t-1}}\lambda\boldsymbol{\pi}_t$, with equality if $q_t^k > 0$. Substituting $\frac{1}{\beta^{t-1}}\lambda\boldsymbol{\pi}_t$ for $\nabla u(\mathbf{q}_t)$ therefore preserves the inequality and gives statement (2):

$$u_s \leq u_t + \tfrac{1}{\beta^{t-1}}\lambda\boldsymbol{\pi}_t'(\mathbf{q}_s - \mathbf{q}_t) \text{ for } \forall s, t = 1, ..., T$$

$$\pi_t^k = \rho_t^k \text{ for any } q_t^k > 0 \text{ for } \forall k = 1, ...K, t = 1, ..., T$$

where $\boldsymbol{\pi}_t \in \mathbb{R}^K_{++}$.

(2) $\implies$ (1) [*Sufficiency*] By the definition of i-rationalisable, we only need to find support prices $\boldsymbol{\pi}_t \in \mathbb{R}^K_{++}$ that satisfy the first-order conditions given by the life cycle model.

Suppose there exist numbers $u_t \in \mathbb{R}$, $\lambda > 0$, $\beta \in (0, 1]$, and vectors $\boldsymbol{\pi}_t \in \mathbb{R}^K_{++}$ such that

$$u_s \leq u_t + \frac{1}{\beta^{t-1}}[\lambda \boldsymbol{\pi}_t'(\mathbf{q}_s - \mathbf{q}_t)] \text{ for } \forall s,t = 1,....,T$$

$$\pi_t^k = \rho_t^k \text{ for any } q_t^k > 0 \text{ for } \forall k = 1,...K, t = 1,...,T$$

Define $\boldsymbol{\mu}_t = \frac{1}{\beta^{t-1}}\boldsymbol{\pi}_t$, so the inequality can be rewritten as

$$u_s \leq u_t + \lambda \boldsymbol{\mu}_t'(\mathbf{q}_s - \mathbf{q}_t) \text{ for } \forall s,t = 1,....,T$$

which is a special case of the standard Afriat inequalities (where $\lambda_t = \lambda$ for $\forall t$). By Afriat's Theorem, the dataset $\mathbf{D'} = \{\boldsymbol{\mu}_t, \mathbf{q}_t\}_{t=1,...,T}$ satisfies GARP and therefore there exists a locally nonsatiated, concave, continuous utility function $u(.)$ that is monotonic in $\mathbf{q}_t$ such that $\nabla u(\mathbf{q}_t) \leq \lambda \boldsymbol{\mu}_t$, with equality when $q_t^k > 0$ or $q_t^k = 0$ and $k \in \kappa_t$.

Either: 1. Suppose that if $q_t^k = 0$, then $k \in \kappa_t$ i.e. all observed zero purchases are due to rationing, so that $\nabla u(\mathbf{q}_t) = \frac{1}{\beta^{t-1}}\lambda \boldsymbol{\pi}_t$ for $t = 1,...,T$.

Given π_t^k and λ, define $\rho_t^k = \pi_t^k - \frac{\gamma_t^k}{\lambda}$ if $q_t^k = 0$, where $0 < \gamma_t^k < \lambda \pi_t^k$, so that $\rho_t^k = \begin{cases} \pi_t^k & q_t^k > 0 \\ \pi_t^k - \frac{\gamma_t^k}{\lambda} & q_t^k = 0 \end{cases}$ and $\boldsymbol{\rho}_t \in \mathbb{R}_{++}^K$. Substituting for π_t^k obtains the first-order conditions for the life cycle model:

$$\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \frac{1}{\beta^{t-1}}\lambda \rho_t^k \text{ if } q_t^k > 0 \text{ ; } t = 1,...,T, k = 1,...,K$$

$$\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \frac{1}{\beta^{t-1}}[\lambda \rho_t^k + \gamma_t^k] \text{ if } q_t^k = 0; \ t = 1,...,T, k = 1,...,K$$

Or: 2. Suppose that if $q_t^k = 0$, then $k \notin \kappa_t$ i.e. all observed zero purchases are due to corner solutions, so that $\nabla u(\mathbf{q}_t) \leq \frac{1}{\beta^{t-1}}\lambda \boldsymbol{\pi}_t$, with equality when $q_t^k > 0$.

Given π_t^k, define $\rho_t^k = \pi_t^k$ for $q_t^k = 0$, so $\rho_t^k = \pi_t^k$ for $\forall q_t^k$. Substituting for π_t^k obtains the first-order conditions for the life cycle model:

$$\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \frac{1}{\beta^{t-1}}\lambda \rho_t^k \text{ if } q_t^k > 0 \text{ ; } t = 1,...,T, k = 1,...,K$$

$$\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} \leq \frac{1}{\beta^{t-1}}\lambda \rho_t^k \text{ if } q_t^k = 0 \text{ ; } t = 1,...,T, k = 1,...,K \ \square$$

The discounted prices ρ_t can be constructed from the solution to the revealed preference test, which are the numbers $\{\pi_t^k\}_{t=1,...,T}^{k=1,...,K}$ and λ that satisfy the inequality in statement (2).

For nonzero purchases and for zero purchases due to corner solutions, the discounted price equals the support price ($\pi_t^k = \rho_t^k$). For zero purchases due to rationing, the discounted price is a function of the support price π_t^k, λ (the marginal utility of discounted total lifetime expenditure), and γ_t^k. γ_t^k represents the marginal cost to utility from rationing, which is unobservable to the researcher but can be any number such that $0 < \gamma_t^k < \lambda\pi_t^k$, so it is possible for a set of values of γ_t^k to rationalise the data for a given solution to the inequalities.

Theorem 1.1 proved that it is possible to rationalise the data by either treating zero demands as if they were all due to rationing or all due to corner solutions, regardless of the actual reason for not purchasing the good. It is also possible to rationalise the data by allowing the zero demands to be due to a mixture of these reasons, by applying the relevant construction of ρ_t^k to each good. One important point shown by the proof is that the revealed preference conditions are exhaustive: the reason for missing prices only matters when constructing discounted prices ρ_t^k from the solution (if it exists), but does not affect what the solution is or whether or not it exists. Therefore, knowing the reason for missing prices would not impose extra restrictions on the data and the revealed preference test would be the same as in the absence of this information.

The inequalities in statement (2) are similar to those for the static case of Crawford and Polisson (2015), showing that their results for missing prices extend to dynamic settings. These inequalities are non-linear in unknowns, but are linear conditional on β. Since the assumed (empirically plausible) range of β is relatively small, the existence of support prices can therefore be determined through a grid search for β over the interval $(0, 1]$ with arbitrarily fine spacing[15], and checking for a solution at each value of β. Thus it is still possible to test if data is consistent with utility maximisation when prices are partially observed, with-

[15] 0.005, for example, or the smallest unit of currency (for practical purposes).

out having to impute missing prices or know whether prices are missing due to rationing or a corner solution.

1.3 Adding storability

1.3.1 Short Memory Habits (SMH) model

Allowing for storability adds more realism to consumer behaviour models, especially if the goods considered are durable (for example, clothing) or if the time between observations is short (for example, scanner data). It may not be necessary for consumers to purchase every type of good in every period if they already have some of the goods stored at home. The Browning (1989) life cycle model from Section 1.2 assumes consumption independence across time periods and so cannot incorporate storability. However, by drawing an analogy between storability and habits, the Short Memory Habits (SMH) model[16] can be used to model storability; that is, the consumer maximises lifetime utility, including the effects of using the storable goods in past and current periods. This section will discuss the analogy between storability and habits, as well as conceptual differences.

In the SMH model, the instantaneous utility function depends both on the current and past consumption (assumed to be a finite number of periods) of K storable goods, with consumption at time t represented by the vector $\mathbf{q}_t$. SMH models are categorised by the number of lags (R) in the utility function. Without loss of generality, assume that all goods are storable[17]. In the context of missing prices, due to storability, the consumer receives utility from having the good in the periods after the time of purchase. For example, zero demand may be observed in some periods for lightbulbs since a lightbulb can be used for a number

[16] A version of the rational intrinsic habits model developed by Ryder and Heal (1973), Spinnewyn (1981), and Becker, Grossman and Murphy (1994) among other authors. Crawford (2010) gives revealed preference conditions for this model under fully observed prices.

[17] Arguably, all goods are storable to some extent. Non-storable goods have no intertemporal links in utility and so can be treated in the same way as in Section 1.2, in terms of support prices and first-order conditions.

of periods (determined by the number of lags included in the SMH model), but is unusable after then. The consumer derives utility from the lightbulb in all periods it is in use.

To focus on the intuition, first consider $R = 1$ (the 1-lag SMH model). The assumption of perfect foresight is maintained, so the consumer's maximisation problem is:

$$\max_{\mathbf{q}_t \in \mathbb{R}_+^K} \sum_{t=1}^{\infty} \beta^{t-1} u(\mathbf{q}_t, \mathbf{q}_{t-1}) \text{ s.t. } \sum_{t=1}^{\infty} \boldsymbol{\rho}_t' \mathbf{q}_t = Y \text{and } q_t^k = 0 \text{ for any } k \in \kappa_t$$

where $\boldsymbol{\rho}_t$ are vectors of discounted prices for storable goods (defined as in Section 1.2), and β, Y, κ_t are also defined as in Section 1.2. As in the life cycle model, the utility function is assumed to be concave, continuous, and strictly increasing in $\mathbf{q_t}$ and $\mathbf{q}_{t-1}$.

The lagrangian is therefore: $\mathcal{L} = \sum_{t=1}^{\infty} \beta^{t-1} u(\mathbf{q}_t, \mathbf{q}_{t-1}) + \lambda(Y - \sum_{t=1}^{\infty} \boldsymbol{\rho}_t' \mathbf{q}_t) + \gamma_t^k (0 - q_t^k)$ and the first-order conditions are:

$$\beta^{t-1} \frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1})}{\partial q_t^k} + \beta^t \frac{\partial u(\mathbf{q}_{t+1}, \mathbf{q}_t)}{\partial q_t^k} = \lambda \rho_t^k \text{ if } q_t^k > 0$$

$$\beta^{t-1} \frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1})}{\partial q_t^k} + \beta^t \frac{\partial u(\mathbf{q}_{t+1}, \mathbf{q}_t)}{\partial q_t^k} \leq \lambda \rho_t^k \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t$$

$$\beta^{t-1} \frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1})}{\partial q_t^k} + \beta^t \frac{\partial u(\mathbf{q}_{t+1}, \mathbf{q}_t)}{\partial q_t^k} = \lambda \rho_t^k + \gamma_t^k \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t$$

The first-order conditions are similar to those of the life cycle model, but include derivatives with respect to next-period's utility function because purchasing the storable good gives both current and future utility. These first-order conditions can be rewritten in a way similar to the life cycle model by defining appropriate discounted shadow prices (as in Spinnewyn (1981) and Crawford (2010)) to represent the discounted willingness to pay:

$$\tilde{\rho}_t^{t,k} \equiv \frac{1}{\lambda} \beta^{t-1} \frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1})}{\partial q_t^k} \text{ for } \forall t \geq 2$$

$$\tilde{\rho}_t^{t+1,k} \equiv \frac{1}{\lambda} \beta^{t} \frac{\partial u(\mathbf{q}_{t+1}, \mathbf{q}_t)}{\partial q_t^k} \text{ for } \forall t \geq 2$$

Here the superscript refers to the time period of the utility function considered (which determines the exponent on β), and the derivative of the utility function is taken with respect to the time period referred to in the subscript.

Substituting in for $\tilde{\rho}_t^{t,k}$ and $\tilde{\rho}_t^{t+1,k}$, the observed discounted price for nonzero demands ($q_t^k > 0$) is equal to the sum of the current and next period's discounted shadow prices:

$$\rho_t^{t,k} = \tilde{\rho}_t^{t,k} + \tilde{\rho}_t^{t+1,k}$$

Unlike habits, which can be detrimental to the consumer in future periods (so $\tilde{\rho}_t^{t+1,k}$ can be negative), storable goods can only be weakly beneficial for the consumer (there is always the option of disposing of a good if its consumption or storage reduces utility, whereas habits cannot be gotten rid of as easily). Therefore, the discounted shadow price of next period's consumption is restricted to be non-negative: $\rho_t^{t+1,k} \geq 0$.

The support prices $\boldsymbol{\pi}_{\mathbf{t}}^{\mathbf{k}}$ can be written in terms of these shadow prices[18]:

$$\boldsymbol{\pi}_t^k = \begin{cases} [\tilde{\rho}_t^{t,k}, \tilde{\rho}_{t-1}^{t,k}]' & \boldsymbol{\chi}_t^k > \mathbf{0} \\ [\tilde{\rho}_t^{t,k} + \frac{\gamma_t^k}{\lambda}, \tilde{\rho}_{t-1}^{t,k} + \frac{\gamma_{t-1}^k}{\lambda}]' & \boldsymbol{\chi}_t^k = \mathbf{0}; k \in \kappa_t \\ \frac{1}{\lambda}[\frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1})}{\partial q_t^k}, \frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1})}{\partial q_{t-1}^k}]' & \boldsymbol{\chi}_t^k = \mathbf{0}; k \notin \kappa_t \end{cases}$$

where $\boldsymbol{\chi}_t^k = [q_t^k, q_{t-1}^k]'$ and $\tilde{\rho}_t^{t,k} = \rho_t^k + \tilde{\rho}_t^{t+1,k}$. If one element of $\boldsymbol{\chi}_t^k$ is nonzero while the other is zero, then the support prices are given by the appropriate element of $\boldsymbol{\pi}_t^k$ as defined above[19]. The support prices for purchased goods can be interpreted as the sum of the discounted willingness-to-pay for current and future consumption of the storable good. Note that unlike the life cycle model, the vector of support prices $\boldsymbol{\pi}_{\mathbf{t}}$ is now of dimension $2K \times 1$ instead of $K \times 1$, due to the presence of lagged consumption ($\mathbf{q}_{t-1}$) in the utility function.

Definition 1.2. A utility function $u(.)$ ($u : \mathbb{R}_+^K \to \mathbb{R}$) **s1-rationalises** the dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ if there exist support prices $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^{2K}$ (where $\boldsymbol{\pi}_t^k = [\tilde{\rho}_t^{t,k}, \tilde{\rho}_{t-1}^{t,k}]'$ if $\boldsymbol{\chi}_t^k > \mathbf{0}$)

[18] At this point it is important to distinguish between the various types of prices used in this chapter, to avoid confusion.

Discounted support prices $\{\pi_t^k\}_{t=1,...,T}^{k=1,...,K}$ are numbers that generate exactly the same demand choices when used in the unrationed maximisation problem as the demand choices made in the rationed maximisation problem. Non-discounted support prices are denoted with a hat ($\hat{\pi}_t^k$).

A **shadow price** $\{\rho_{t-i}^{t+i,k}\}_{i=0,1}^{k=1,...,k}$ is the value of the lagrange multiplier at the optimal solution, which can be interpreted as the discounted willingness-to-pay.

The **discounted price** ρ_t^k is the period-1 price of good k that was available for sale in a future period t, adjusted for the interest rate i.e. $\rho_t^k = \frac{p_t^k}{\prod_{s=2}^{s=t}(1+i_s)}$, where i_s is the interest rate in period s and p_t^k is the price of good k in period $t \in \{2, ...T\}$.

[19] For example, if $\boldsymbol{\chi}_t^k = [q_t^k, 0]'$, and the zero demand was due to rationing, then $\boldsymbol{\pi}_t^k = [\tilde{\rho}_t^{t,k}, \tilde{\rho}_{t-1}^{t,k} + \frac{\gamma_{t-1}^k}{\lambda}]'$.

and $\beta \in (0,1]$ such that $u(\boldsymbol{\chi}_t) \geq u(\boldsymbol{\chi})$ for any $\boldsymbol{\chi} \in \{\chi \in \mathbb{R}^k_+ : \boldsymbol{\pi}_t^{'}\boldsymbol{\chi} \leq \boldsymbol{\pi}_t^{'}\boldsymbol{\chi}_t\}$ for $\forall\ t \in \{1,...,T\}$, where $\boldsymbol{\chi}^k_t = [q^k_t, q^k_{t-1}]'$.

Applying Definition 1.2 in this context, the dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ (where $q^k_t = 0$ for some $t \in \{2,...,T\}$, $k \in \{1,...,K\}$) is therefore rationalisable by the one-lag SMH model (s1-rationalisable) if there exists a well-behaved utility function $u(.)$ and numbers $\lambda > 0$, $\gamma^k_t > 0$, $\beta \in (0,1]$ such that the FOCs of optimising behaviour are satisfied, i.e.

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1})}{\partial q^k_t} + \beta^t \frac{\partial u(\mathbf{q}_{t+1},\mathbf{q}_t)}{\partial q^k_t} = \lambda \rho^k_t \text{ if } q^k_t > 0$$

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1})}{\partial q^k_t} + \beta^t \frac{\partial u(\mathbf{q}_{t+1},\mathbf{q}_t)}{\partial q^k_t} \leq \lambda \rho^k_t \text{ if } q^k_t = 0 \text{ and } k \notin \kappa_t$$

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1})}{\partial q^k_t} + \beta^t \frac{\partial u(\mathbf{q}_{t+1},\mathbf{q}_t)}{\partial q^k_t} = \lambda \rho^k_t + \gamma^k_t \text{ if } q^k_t = 0 \text{ and } k \in \kappa_t.$$

As in Section 1.2, the observed choices are consistent with utility maximisation if we can find a utility function and support prices (appropriate shadow prices for purchased goods, and reservation/virtual prices for goods not purchased) such that the first-order conditions are satisfied i.e. solving the first-order conditions and substituting in the support prices will result in the observed demands. Theorem 1.2 establishes the necessary and sufficient conditions for data to be consistent with the SMH model, even when some prices are missing.

Theorem 1.2 (1-lag SMH model). The following statements are equivalent:

(1) The data $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ is s1-rationalisable.

(2) There exist numbers $u_t \in \mathbb{R}$, $\lambda > 0$, $\beta \in (0,1]$, $\{\tilde{\rho}^{t,k}_{t-i}\}^{k=1,...,K}_{i=0,1} > 0$, $\{\tilde{\rho}^{t+1,k}_t\}^{k=1,...,K} \geq 0$ and vectors $\boldsymbol{\pi}_t \in \mathbb{R}^{2K}_{++}$ such that

$$u_s \leq u_t + \lambda \boldsymbol{\pi}_t^{'}(\boldsymbol{\chi}_s - \boldsymbol{\chi}_t) \text{ for } \forall s,t = 2,....,T$$

$$\rho^k_t = \tilde{\rho}^{t,k}_t + \tilde{\rho}^{t+1,k}_t \text{ for any } q^k_t > 0,\ k = 1,...,K,\ t \in \{2,...,T\}$$

$$\boldsymbol{\pi}^k_t = \frac{1}{\beta^{t-1}}[\tilde{\rho}^{t,k}_t, \tilde{\rho}^{t,k}_{t-1}]' \text{ for any } \boldsymbol{\chi}^k_t > 0,\ k = 1,...,K, t \in \{2,...,T\}$$

where $\boldsymbol{\chi}^k_t = [q^k_t, q^k_{t-1}]'$.

Proof:

(1) $\implies$ (2) [*Necessity*] The dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ is rationalisable by the one-lag SMH model if there exists a locally non-satiated, continuous, concave, and differentiable utility function $u(.)$ and numbers $\lambda > 0$, $\gamma_t^k > 0$, $\beta \in (0, 1]$ such that

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1})}{\partial q_t^k} + \beta^t\frac{\partial u(\mathbf{q}_{t+1},\mathbf{q}_t)}{\partial q_t^k} = \lambda\rho_t^k \text{ if } q_t^k > 0$$

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1})}{\partial q_t^k} + \beta^t\frac{\partial u(\mathbf{q}_{t+1},\mathbf{q}_t)}{\partial q_t^k} \leq \lambda\rho_t^k \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t$$

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1})}{\partial q_t^k} + \beta^t\frac{\partial u(\mathbf{q}_{t+1},\mathbf{q}_t)}{\partial q_t^k} = \lambda\rho_t^k + \gamma_t^k \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t$$

These first-order conditions can be rearranged in terms of $\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1})}{\partial q_t^k}$:

$$\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1})}{\partial q_t^k} = \frac{1}{\beta^{t-1}}\lambda\tilde{\rho}_t^{t,k} \text{ if } q_t^k > 0 \text{ ; } t = 2, ..., T, k = 1, ..., K$$

$$\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1})}{\partial q_t^k} \leq \frac{1}{\beta^{t-1}}\lambda\tilde{\rho}_t^{t,k} \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t \text{ ; } t = 2, ..., T, k = 1, ..., K$$

$$\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1})}{\partial q_t^k} = \frac{1}{\beta^{t-1}}\lambda[\tilde{\rho}_t^{t,k} + \frac{\gamma_k^k}{\lambda}] \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t\text{; } t = 2, ..., T, k = 1, ..., K$$

where $\tilde{\rho}_t^{t,k} = \rho_t^k - \tilde{\rho}_t^{t+1,k}$.

By concavity of $u(.)$,

$$u_s \leq u_t + \nabla u(\boldsymbol{\chi}_t)'(\boldsymbol{\chi}_s - \boldsymbol{\chi}_t) \text{ for } \forall s, t = 2, ..., T$$

where $\boldsymbol{\chi}_t^k = [q_t^k, q_{t-1}^k]'$.

For $t = 2, ..., T, k = 1, ..., K$, define $\tilde{\pi}_t^k = \frac{1}{\beta^{t-1}}\tilde{\rho}_t^{t,k}$ if $q_t^k > 0$ or $q_t^k = 0$ and $k \notin \kappa_t$, and $\tilde{\pi}_t^k = \frac{1}{\beta^{t-1}}(\tilde{\rho}_t^{t,k} + \frac{\gamma_t^k}{\lambda})$ if $q_t^k = 0$ and $k \in \kappa_t$. Define $\boldsymbol{\pi}_t^k = [\tilde{\pi}_t^k, \tilde{\pi}_{t-1}^k]'$. Note that $\boldsymbol{\pi}_t^k \in \mathbb{R}_{++}^2$ since $\tilde{\rho}_t^{t,k}$, γ_t^k, and $\lambda > 0$, so $\nabla u(\boldsymbol{\chi}_t) \leq \lambda\boldsymbol{\pi}_t$, with equality if $q_t^k > 0$. Substituting $\lambda\boldsymbol{\pi}_t$ for $\nabla u(\mathbf{q}_t)$ therefore preserves the inequality and gives statement (2):

$$u_s \leq u_t + \lambda\boldsymbol{\pi}_t'(\boldsymbol{\chi}_s - \boldsymbol{\chi}_t) \text{ for } \forall s, t = 2, ..., T$$

$$\rho_t^k = \tilde{\rho}_t^{t,k} + \tilde{\rho}_t^{t+1,k} \text{ for any } q_t^k > 0,\ k = 1, ..., K,\ t \in \{2, ..., T\}$$

$$\boldsymbol{\pi}_t^k = \frac{1}{\beta^{t-1}}[\tilde{\rho}_t^{t,k}, \tilde{\rho}_{t-1}^{t,k}]' \text{ for any } \boldsymbol{\chi}_t^k > 0,\ k = 1, ..., K, t \in \{2, ..., T\}$$

where $\boldsymbol{\pi}_t \in \mathbb{R}^{2K}_{++}$.

(2) $\Longrightarrow$ (1) [*Sufficiency*] By the definition of s1-rationalisable, we only need to find support prices $\boldsymbol{\pi}_t \in \mathbb{R}^{2K}_{++}$ that satisfy the first-order conditions given by the one-lag SMH model.

Suppose there exist numbers $u_t \in \mathbb{R}$, $\lambda > 0$, $\beta \in (0,1]$, $\tilde{\rho}^{t,k}_{t-1} > 0, \tilde{\rho}^{t,k}_t > 0$, $\tilde{\rho}^{t+1,k}_t \neq 0$, and vectors $\boldsymbol{\pi}_t \in \mathbb{R}^K_{++}$ such that

$$u_s \leq u_t + \lambda \boldsymbol{\pi}'_t(\boldsymbol{\chi}_s - \boldsymbol{\chi}_t) \text{ for } \forall s,t = 2, ..., T$$

$$\rho^k_t = \tilde{\rho}^{t,k}_t + \tilde{\rho}^{t+1,k}_t \text{ for any } q^k_t > 0,\ k = 1, ..., K,\ t \in \{2, ..., T\}$$

$$\boldsymbol{\pi}^k_t = \tfrac{1}{\beta^{t-1}}[\tilde{\rho}^{t,k}_t, \tilde{\rho}^{t,k}_{t-1}]' \text{ for any } \boldsymbol{\chi}^k_t > 0,\ k = 1, ..., K, t \in \{2, ..., T\}$$

The inequalities $\{u_s \leq u_t + \lambda \boldsymbol{\pi}'_t(\boldsymbol{\chi}_s - \boldsymbol{\chi}_t)\}_{s,t=1,...,T}$ are a special case of the standard Afriat inequalities (where $\lambda_t = \lambda$ for $\forall t$). By Afriat's Theorem, the dataset $\mathbf{D}'' = \{\boldsymbol{\pi}_t, \boldsymbol{\chi}_t\}_{t=1,...,T}$ satisfies GARP and therefore there exists a locally nonsatiated, concave, continuous utility function $u(.)$ that is monotonic in $\boldsymbol{\chi}_t$ such that $\nabla u(\boldsymbol{\chi}_t) \leq \lambda \boldsymbol{\pi}_t$, with equality when $q^k_t > 0$ or $q^k_t = 0$ and $k \in \kappa_t$.

<u>Either:</u> 1. Suppose that if $q^k_t = 0$, then $k \in \kappa_t$ i.e. all observed zero purchases are due to rationing, so that $\nabla u(\boldsymbol{\chi}_t) = \lambda \boldsymbol{\pi}_t$ for $t = 2, ..., T$.

Consider the first row of $\nabla_k u(\boldsymbol{\chi}_t)$: $\frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1})}{\partial q^k_t} = \lambda \tilde{\pi}^k_t$. Given $\tilde{\pi}^k_t$, $\tilde{\rho}^{t+1,k}_t$, β, and λ, define $\tilde{\rho}^k_t = \beta^{t-1}\tilde{\pi}^k_t + \tilde{\rho}^{t+1,k}_t - \frac{\gamma^k_t}{\lambda}$ if $q^k_t = 0$, where $0 < \gamma^k_t < \lambda[\beta^{t-1}\tilde{\pi}^k_t + \tilde{\rho}^{t+1,k}_t]$, so that $\rho^k_t = \begin{cases} \beta^{t-1}\tilde{\pi}^k_t + \tilde{\rho}^{t+1,k}_t & q^k_t > 0 \\ \beta^{t-1}\tilde{\pi}^k_t + \tilde{\rho}^{t+1,k}_t - \frac{\gamma^k_t}{\lambda} & q^k_t = 0 \end{cases}$ and $\boldsymbol{\rho}_t \in \mathbb{R}^{2K}_{++}$. Substituting for $\tilde{\pi}^k_t$ obtains the first-order conditions for the one-lag SMH model:

$$\tfrac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1})}{\partial q^k_t} = \tfrac{1}{\beta^{t-1}}\lambda \tilde{\rho}^{t,k}_t \text{ if } q^k_t > 0 \ ;\ t = 2, ..., T, k = 1, ..., K$$

$$\tfrac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1})}{\partial q^k_t} = \tfrac{1}{\beta^{t-1}}\lambda[\tilde{\rho}^{t,k}_t + \tfrac{\gamma^k_k}{\lambda}] \text{ if } q^k_t = 0 \text{ and } k \in \kappa_t;\ t = 2, ..., T, k = 1, ..., K$$

where $\tilde{\rho}^{t,k}_t = \rho^k_t - \tilde{\rho}^{t+1,k}_t$.

Or: 2. Suppose that if $q_t^k = 0$, then $k \notin \kappa_t$ i.e. all observed zero purchases are due to corner solutions, so that $\nabla u(\boldsymbol{\chi}_t) \leq \lambda \boldsymbol{\pi}_t$, with equality when $q_t^k > 0$.

Consider the first row of $\nabla_k u(\boldsymbol{\chi}_t)$: $\frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1})}{\partial q_t^k} < \lambda \tilde{\pi}_t^k$, with equality when $q_t^k > 0$.

Given $\tilde{\pi}_t^k$, define $\rho_t^k = \beta^{t-1}\tilde{\pi}_t^k + \tilde{\rho}_t^{t+1,k}$ for $q_t^k = 0$, so $\rho_t^k = \beta^{t-1}\tilde{\pi}_t^k + \tilde{\rho}_t^{t+1,k}$ for $\forall q_t^k$. Substituting for $\tilde{\pi}_t^k$ obtains the first-order conditions for the one-lag SMH model:

$$\frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1})}{\partial q_t^k} = \frac{1}{\beta^{t-1}} \lambda \tilde{\rho}_t^{t,k} \text{ if } q_t^k > 0 \ ; \ t = 2, ..., T, k = 1, ..., K$$

$$\frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1})}{\partial q_t^k} \leq \frac{1}{\beta^{t-1}} \lambda \tilde{\rho}_t^{t,k} \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t \ ; \ t = 2, ..., T, k = 1, ..., K$$

where $\tilde{\rho}_t^{t,k} = \rho_t^k - \tilde{\rho}_t^{t+1,k}$. $\square$

The revealed preference conditions with partially observed prices are a generalisation of those in Crawford (2010), with the conditions relating shadow prices to support prices and discounted prices now only required to hold for nonzero demands. Now that utility depends on consumption in the previous period, there are $T - 1$ instead of T distinct inequalities. As in the life cycle model, these inequalities are non-linear in unknowns, but are linear conditional on β. Since the assumed range of β is relatively small, the existence of support prices can again be determined by using a grid search for β within the range $(0, 1]$ with arbitrarily fine spacing, and running a linear programming problem for each value of β. The revealed preference conditions are also exhaustive for the same reason as in Section 1.2, so would not change even if the reason(s) for missing prices were known.

These results for the 1-lag SMH model can be generalised to a SMH model with $R > 1$ lags, where the consumer's instantaneous utility function is now $u(\mathbf{q}_t, \mathbf{q}_{t-1}, ..., \mathbf{q}_{t-R})$. The first order conditions are:

$$\beta^{t-1} \frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1}, ..., \mathbf{q}_{t-R})}{\partial q_t^k} + \sum_{r=1}^R \beta^{t+r-1} \frac{\partial u(\mathbf{q}_{t+r-1}, \mathbf{q}_{t+r-2}, ..., \mathbf{q}_{t+r-R})}{\partial q_t^k} = \lambda \rho_t^k \text{ if } q_t^k > 0$$

$$\beta^{t-1} \frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1}, ..., \mathbf{q}_{t-R})}{\partial q_t^k} + \sum_{r=1}^R \beta^{t+r-1} \frac{\partial u(\mathbf{q}_{t+r-1}, \mathbf{q}_{t+r-2}, ..., \mathbf{q}_{t+r-R})}{\partial q_t^k} \leq \lambda \rho_t^k \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t$$

$$\beta^{t-1} \frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1}, ..., \mathbf{q}_{t-R})}{\partial q_t^k} + \sum_{r=1}^R \beta^{t+r-1} \frac{\partial u(\mathbf{q}_{t+r-1}, \mathbf{q}_{t+r-2}, ..., \mathbf{q}_{t+r-R})}{\partial q_t^k} = \lambda \rho_t^k + \gamma_t^k \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t$$

and the support prices $\boldsymbol{\pi}_t^k$ are:

$$\boldsymbol{\pi}_t^k = \begin{cases} [\tilde{\rho}_t^{t,k}, \tilde{\rho}_{t-1}^{t,k}, ..., \tilde{\rho}_{t-R}^{t,k}]' & \boldsymbol{\chi}_t^k > \mathbf{0} \\ [\tilde{\rho}_t^{t,k} + \frac{\gamma_t^k}{\lambda}, \tilde{\rho}_{t-1}^{t,k} + \frac{\gamma_{t-1}^k}{\lambda}, ..., \tilde{\rho}_{t-R}^{t,k} + \frac{\gamma_{t-R}^k}{\lambda}]' & \boldsymbol{\chi}_t^k = \mathbf{0}; k \in \kappa_t \\ \frac{1}{\lambda}[\frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1}, ..., \mathbf{q}_{t-R})}{\partial q_t^k}, ..., \frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1}, ..., \mathbf{q}_{t-R})}{\partial q_{t-R}^k}]' & \boldsymbol{\chi}_t^k = \mathbf{0}; k \notin \kappa_t \end{cases}$$

where $\{\tilde{\rho}_{t-r}^{t,k}\}_{r=0,..,R}$ are discounted shadow prices defined as in the one-lag SMH model, $\tilde{\rho}_t^{t,k} = \rho_t^k - \sum_{r=1}^R \tilde{\rho}_t^{t+r,k}$, and $\boldsymbol{\chi}_t^k = [q_t^k, q_{t-1}^k, ..., q_{t-R}^k]'$. Again, the support prices for a general vector $\boldsymbol{\chi}_t^k$ are given by the appropriate elements of $\boldsymbol{\pi}_t^k$ as defined above. Note that now the vector of support prices $\boldsymbol{\pi}_t^k$ is of dimension $K(R+1) \times 1$. There is also a similar relationship between discounted shadow prices and observed discounted prices as in the 1-lag SMH model, where the observed discounted price is equal to the sum of discounted willingness to pay over the current and R future periods:

$$\rho_t^k = \sum_{i=0}^R \tilde{\rho}_t^{t+r,k} \text{ for } q_t^k > 0$$

In other words, the discounted shadow price of current consumption for nonzero demands ($q_t^k > 0$) is equal to the observed discounted price, taking into account the welfare effects from using the good in all R future periods. As in the 1-lag SMH model, storability must be weakly beneficial for the consumer, so the discounted shadow prices of consumption in future periods $\{\tilde{\rho}_t^{t+r,k}\}^{r=1,...R}$ should be non-negative.

Substituting $\tilde{\rho}_t^{t,k}$ recursively into the first order condition for purchased goods gives

$$\tilde{\rho}_t^{t,k} = \lambda \rho_t^k - \sum_{r=1}^R \tilde{\rho}_t^{t+r,k}$$

which shows the intuition that in equilibrium, the marginal benefit of an increase in consumption of q_t^k should equal its marginal cost (Demunyck & Verriest, 2013).

Definition 1.3. A utility function $u(.)$ $(u : \mathbb{R}_+^K \to \mathbb{R})$ **sr-rationalises** the dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ if there exist support prices $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^{K(R+1)}$ (where $\boldsymbol{\pi}_t^k = [\tilde{\rho}_t^{t,k}, \tilde{\rho}_{t-1}^{t,k}, ..., \tilde{\rho}_{t-R}^{t,k}]'$

if $\boldsymbol{\chi}_t^k > \mathbf{0}$) and $\beta \in (0,1]$ such that $u(\boldsymbol{\chi}_t) \geq u(\boldsymbol{\chi})$ for any $\boldsymbol{\chi} \in \{\chi \in \mathbb{R}_+^k : \boldsymbol{\pi}_t^{'}\boldsymbol{\chi} \leq \boldsymbol{\pi}_t^{'}\boldsymbol{\chi}_t\}$ for $\forall$ $t \in \{1, ..., T\}$, where $\boldsymbol{\chi}_t^k = [q_t^k, q_{t-1}^k, ..., q_{t-R}^k]'$.

By Definition 1.3, the dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ (where $q_t^k = 0$ for some $t \in \{R+1, ..., T\}$, $k \in \{1, ..., K\}$) is rationalisable by the R-lag SMH model (sr-rationalisable) if there exists a locally non-satiated, continuous, concave, and differentiable utility function $u(.)$ and numbers $\lambda > 0$, $\gamma_t^k > 0$, $\beta \in \,]0,1]$ such that the FOCs of optimising behaviour are satisfied, i.e.

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1},...,\mathbf{q}_{t-R})}{\partial q_t^k} + \sum_{r=1}^{R} \beta^{t+r-1}\frac{\partial u(\mathbf{q}_{t+r-1},\mathbf{q}_{t+r-2},...,\mathbf{q}_{t+r-R})}{\partial q_t^k} = \lambda \rho_t^k \text{ if } q_t^k > 0$$

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1},...,\mathbf{q}_{t-R})}{\partial q_t^k} + \sum_{r=1}^{R} \beta^{t+r-1}\frac{\partial u(\mathbf{q}_{t+r-1},\mathbf{q}_{t+r-2},...,\mathbf{q}_{t+r-R})}{\partial q_t^k} \leq \lambda \rho_t^k \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t$$

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1},...,\mathbf{q}_{t-R})}{\partial q_t^k} + \sum_{r=1}^{R} \beta^{t+r-1}\frac{\partial u(\mathbf{q}_{t+r-1},\mathbf{q}_{t+r-2},...,\mathbf{q}_{t+r-R})}{\partial q_t^k} = \lambda \rho_t^k + \gamma_t^k \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t$$

Theorem 1.3 states the necessary and sufficient conditions on observables for the data to be consistent with the R-lag SMH model.

Theorem 1.3 (R-lag SMH model). The following statements are equivalent:

(1) The data $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ is sr-rationalisable.

(2) There exist numbers $u_t \in \mathbb{R}$, $\lambda > 0$, $\beta \in \,]0,1]$, $\{\tilde{\rho}_{t-i}^{t,k}\}_{i=0,...,R}^{k=1,...,K} > 0$, $\{\tilde{\rho}_t^{t+i,k}\}^{i=0,...,R;k=1,...,K} \neq 0$, and vectors $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^{K(R+1)}$ such that

$$u_s \leq u_t + \frac{1}{\beta^{t-1}}[\boldsymbol{\pi}_t^{'}(\boldsymbol{\chi}_s - \boldsymbol{\chi}_t)] \text{ for } \forall s,t = R+1, ..., T$$

$$\rho_t^k = \sum_{r=0}^{R} \tilde{\rho}_t^{t+r,k} \text{ for any } q_t^k > 0,\ k = 1, ..., K, t \in \{R+1, ..., T\}$$

$$\boldsymbol{\pi}_t^k = [\tilde{\rho}_t^{t,k}, \tilde{\rho}_{t-1}^{t,k}, ..., \tilde{\rho}_{t-R}^{t,k}]' \text{ for any } \boldsymbol{\chi}_t^k > 0,\ k = 1, ..., K, t \in \{R+1, ..., T\}$$

where $\boldsymbol{\chi}_t^k = [q_t^k, q_{t-1}^k, ..., q_{t-R}^k]'$.

Proof: See Appendix 1A.

These revealed preference conditions are a generalisation of those given in Crawford (2010): for purchased goods the conditions are identical, and for goods not purchased there must be numbers (support prices) that make these zero demands rational when taking observed purchases into account. The revealed preference conditions for the R-lag model only differ according to the number of discounted shadow prices present (i.e. number of future periods in which the consumer can use the good). Note that for the R-lag SMH model, there are $T - R$ distinct inequalities because utility depends on consumption over the past R periods, so the revealed preference test omits the first R time periods. The number of time periods observed can thus substantially restrict the number of lags it is possible to test for, especially if only a few time periods are observed. This issue is less of a problem in the next habits model considered.

1.3.2 Habits as Durables (HAD) model

The HAD model, first introduced by Becker, Grossman, and Murphy (1988) and developed by Demuynck and Verriest (2013), is another way to model storability. Consumption of a storable good (q_t) is likened to "investing" in a stock of this good, S_t, which is also part of the instantaneous utility function. The consumer thus gets utility from current consumption, but also from past consumption through the stock variable. The stock variable S_t represents the cumulative stock of that item the consumer has stored at home, adjusted for 'depreciation'. As in standard investment problems, this stock depreciates at a fixed rate, $\delta \in (0, 1]$, and grows according to current and past consumption of the storable good. For example, zero demand may be observed for a bag of apples in some periods since the consumer already has stored a large stock of apples, so the total stock of apples is whatever amount was purchased last period, minus the amount of apples that have gone bad. δ, the proportion of the stock that is left for next period, can be interpreted as how storable the good is. Therefore current consumption q_t can potentially affect welfare in all future periods, unlike the SMH model, where current consumption affects future utility only for a finite number of periods.

First, for simplicity, assume there is only one storable good instead of K storable goods, so the instantaneous utility function is $u(q_t, S_t)$, where q_t and S_t are scalars, and assume that discounted lifetime wealth Y is exogenous, with no saving or borrowing[20]. As in previous sections, the instantaneous utility function is assumed to be concave, continuous, and strictly increasing in q_t. Since storability is weakly beneficial for the consumer, $u(q_t, S_t)$ is increasing in S_t. As in Demuynck and Verriest (2013), for $t \geq 2$ the stock evolves according to the following equation:

$$S_t = (1 - \delta)S_{t-1} + q_{t-1}$$

The assumption of perfect foresight is also maintained, so the consumer's maximisation problem is:

$$\max_{q_t \in \mathbb{R}_+} \sum_{t=1}^{\infty} \beta^{t-1} u(q_t, S_t) \text{ s.t. } \sum_{t=1}^{\infty} \rho_t q_t = Y$$

$$S_t = (1 - \delta)S_{t-1} + q_{t-1} \text{ for } \forall t \geq 2,\ S_1 = \bar{S}_1$$

$$q_t = 0 \text{ for any } t \in \sigma$$

where ρ_t are discounted prices of the storable good, $\bar{S}_1$ is the initial stock of the good, and $\sigma \subseteq \{1, ..., T\}$ is the subset of time periods in which the good q is unavailable.

The lagrangian is therefore: $\mathcal{L} = \sum_{t=1}^{\infty} \beta^{t-1} u(q_t, S_t) + \lambda(Y - \sum_{t=1}^{\infty} \rho_t q_t) + \theta_t(S_t - (1 - \delta)S_{t-1} + q_{t-1}) + \gamma_t(0 - q_t)$ and the first-order conditions are:

$$\beta^{t-1} \frac{\partial u(q_t, S_t)}{\partial q_t} = \lambda \rho_t + \theta_t \text{ if } q_t > 0$$

$$\beta^{t-1} \frac{\partial u(q_t, S_t)}{\partial q_t} \leq \lambda \rho_t + \theta_t \text{ if } q_t = 0 \text{ and } t \notin \sigma$$

$$\beta^{t-1} \frac{\partial u(q_t, S_t)}{\partial q_t} = \lambda \rho_t + \theta_t + \gamma_t \text{ if } q_t = 0 \text{ and } t \in \sigma$$

$$\beta^{t-1} \frac{\partial u(q_t, S_t)}{\partial S_t} = (1 - \delta)\theta_t - \theta_{t-1}$$

[20]Demuynck and Verriest (2013) included saving and borrowing in their revealed preference analysis of the HAD model, and also allowed for imperfect capital markets (borrowing constraints). The additional revealed preference conditions implied by these conditions do not depend on whether prices are partially observed or not, so for simplicity the assumption of perfect capital markets is maintained here.

Unlike the SMH model, where the intertemporal link comes directly from having past consumption in the utility function, in the HAD model the intertemporal link comes from $\theta_t < 0$, which is the marginal increase in lifetime utility from a marginal increase in next period's stock (S_{t+1}). The consumer accounts for the effects of increasing the stock on utility when choosing optimal consumption of the storable good. The first-order condition for the stock variable shows that increasing the stock of the good has direct benefits this period (from increasing S_t) but also indirect benefits next period since increasing S_t also increases S_{t+1}.

As in Demuynck & Verriest (2013), these first-order conditions can be rewritten by defining the discounted shadow prices (discounted willingness to pay) of the storable good and the stock variable as:

$$\rho_t^q \equiv \beta^{t-1} \frac{\partial u(q_t, S_t)}{\partial q_t} \text{ for } \forall t \geq 1$$

$$\rho_t^S \equiv \beta^{t-1} \frac{\partial u(q_t, S_t)}{\partial S_t} \text{ for } \forall t \geq 1$$

$\rho_t^q \geq 0$ represents the discounted gain in marginal utility from an increase in consumption of the storable good, while $\rho_t^S \geq 0$ represents the discounted marginal utility from an increase in the stock variable (non-negative since storability is weakly beneficial). The FOC for the stock variable can be rewritten in terms of these shadow prices:

$$\rho_t^S = (1-\delta)[\rho_t^q - \lambda \rho_t] - [\rho_{t-1}^q - \lambda \rho_{t-1}]$$

Substituting for ρ_t^S recursively into the first order condition for purchased goods gives

$$\rho_t^q + \sum_{i=1}^{\infty} (1-\delta)^{i-1} \rho_{t+i}^S = \lambda \rho_t$$

which shows again that in equilibrium, the marginal benefit of consumption must equal the marginal cost (Demuynck & Verriest, 2013). The marginal benefit in this case is the sum of both current benefits (marginal utility from using the good this period) and future benefits from storing the good (due to the increase in the future stock variable).

The support prices π_t^k are:

$$\pi_t^k = \begin{cases} \lambda\rho_t + \theta_t & q_t^k > 0 \\ \lambda\rho_t + \theta_t + \gamma_t & q_t^k = 0; t \in \sigma \\ \frac{1}{\lambda}[\frac{\partial u(q_t,S_t)}{\partial q_t} - \theta_t] & q_t^k = 0; t \notin \sigma \end{cases}$$

Definition 1.4. A utility function $u(.)$ ($u : \mathbb{R}_+^K \to \mathbb{R}$) **h-rationalises** the dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ if there exist support prices $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^K$ (where $\pi_t^k = \lambda\rho_t + \theta_t$ if $q_t^k > 0$) and $\beta, \delta \in (0,1]$ such that for $\forall\ t \in \{1,...,T\}$, $u(\mathbf{q}_t, \mathbf{S}_t) \geq u(\mathbf{q}, \mathbf{S}_t)$ for any $\mathbf{q} \in \{\mathbf{q} \in \mathbb{R}_+^k : \boldsymbol{\pi}_t^{'}\mathbf{q} \leq \boldsymbol{\pi}_t^{'}\mathbf{q}_t\}$.

By Definition 1.4, the dataset $\mathbf{D} = \{i_t, p_t, q_t\}_{t=1,...,T}$ (where $q_t^k = 0$ for some $t \in \{1,...,T\}$, $k \in \{1,...,K\}$) is therefore rationalisable by the HAD model (h-rationalisable) if there exists a locally non-satiated, continuous, concave, and differentiable utility function $u(.)$, and numbers $\lambda > 0$, $\theta_t \geq 0$, $\gamma_t > 0, S_t > 0$, and $\beta, \delta \in (0,1]$ such that the FOCs of optimising behaviour are satisfied i.e.

$$\beta^{t-1}\frac{\partial u(q_t,S_t)}{\partial q_t} = \lambda\rho_t + \theta_t \text{ if } q_t > 0$$

$$\beta^{t-1}\frac{\partial u(q_t,S_t)}{\partial q_t} \leq \lambda\rho_t + \theta_t \text{ if } q_t = 0 \text{ and } t \notin \sigma$$

$$\beta^{t-1}\frac{\partial u(q_t,S_t)}{\partial q_t} = \lambda\rho_t + \theta_t + \gamma_t \text{ if } q_t = 0 \text{ and } t \in \sigma$$

$$\beta^{t-1}\frac{\partial u(q_t,S_t)}{\partial S_t} = (1-\delta)\theta_t - \theta_{t-1}$$

Theorem 1.4 states the necessary and sufficient conditions on observables for the data to be rationalisable by the HAD model.

Theorem 1.4 (1-good HAD model). The following statements are equivalent:

(1) The data $\mathbf{D} = \{i_t, p_t, q_t\}_{t=1,...,T}$ is h-rationalisable.

(2) There exist numbers $u_t, S_t > 0$,$\lambda > 0$, $\rho_t^q > 0$, $\rho_t^S \leq 0$, δ, $\beta \in (0,1]$ and $\pi_t > 0$ such that

$$u_s \leq u_t + \frac{1}{\beta^{t-1}}[\pi_t(q_s - q_t) + \rho_t^S(S_s - S_t)] \text{ for } \forall s,t = 2,...,T \qquad [1]$$

$$\pi_t = \rho_t^q \text{ if } q_t > 0, t = 2,...,T \qquad [2]$$

$$\rho_t^S = (1-\delta)[\rho_t^q - \lambda\rho_t] - [\rho_{t-1}^q - \lambda\rho_{t-1}] \text{ if } q_t > 0,\ t = 2, ..., T \qquad [3]$$

$$S_t = (1-\delta)S_{t-1} + q_{t-1} \text{ for } t = 2, ..., T \qquad [4]$$

Proof:

(1) $\implies$ (2) [*Necessity*] The dataset $\mathbf{D} = \{i_t, p_t, q_t\}_{t=1,...,T}$ is rationalisable by the HAD model if there exists a locally non-satiated, continuous, concave, and differentiable utility function $u(.)$ and numbers $\lambda > 0$, $\theta_t > 0$, $\gamma_t \in \mathbb{R}$, $\delta \in (0,1]$, and $\beta \in (0,1]$ such that

$$\tfrac{\partial u(q_t,S_t)}{\partial q_t} = \tfrac{1}{\beta^{t-1}}[\lambda\rho_t + \theta_t] \text{ if } q_t > 0 \ ;\ t = 2, .., T$$

$$\tfrac{\partial u(q_t,S_t)}{\partial q_t} \leq \tfrac{1}{\beta^{t-1}}[\lambda\rho_t + \theta_t] \text{ if } q_t = 0 \text{ and } t \notin \sigma;\ t = 2, ..., T$$

$$\tfrac{\partial u(q_t,S_t)}{\partial q_t} = \tfrac{1}{\beta^{t-1}}[\lambda\rho_t + \theta_t + \gamma_t] \text{ if } q_t = 0 \text{ and } t \in \sigma;\ t = 2, ..., T$$

$$\tfrac{\partial u(q_t,S_t)}{\partial S_t} = \tfrac{1}{\beta^{t-1}}[(1-\delta)\theta_t - \theta_{t-1}] \ ;\ t = 2, ..., T$$

Conditions [3] is obtained by substituting for θ_t and θ_{t-1} into the first-order condition for stock variable, and using the definitions of the discounted shadow prices ρ_t^S and ρ_t^q. Condition [4] is the equation of motion for the stock variable.

By concavity of $u(.)$,

$$u_s \leq u_t + \tfrac{\partial u(q_t,S_t)}{\partial q_t}(q_s - q_t) + \tfrac{\partial u(q_t,S_t)}{\partial S_t}(S_s - S_t) \text{ for } \forall s, t = 2, ..., T$$

For $t = 1, ..., T$, define $\pi_t = \lambda\rho_t + \theta_t$ if $q_t > 0$ or $q_t = 0$ and $t \notin \sigma$, and $\pi_t = \lambda\rho_t + \theta_t + \gamma_t$ if $q_t = 0$ and $t \notin \sigma$. Note that $\pi_t = \beta^{t-1}\frac{\partial u(q_t,S_t)}{\partial q_t} = \rho_t^q$ if $q_t > 0$, so condition [2] holds. $\pi_t > 0$ since $\theta_t \geq 0$, and ρ_t, γ_t, $\lambda > 0$, so $\frac{\partial u(q_t,S_t)}{\partial q_t} \leq \frac{1}{\beta^{t-1}}\pi_t$, with equality if $q_t > 0$. Also define $\rho_t^S = (1-\delta)\theta_t - \theta_{t-1}$ for $t = 2, ..., T$, so $\frac{\partial u(q_t,S_t)}{\partial S_t} = \frac{1}{\beta^{t-1}}\rho_t^S$.

Substituting $\frac{1}{\beta^{t-1}}\pi_t$ for $\frac{\partial u(q_t,S_t)}{\partial q_t}$ and $\frac{1}{\beta^{t-1}}\rho_t$ for $\frac{\partial u(q_t,S_t)}{\partial S_t}$ therefore preserves the inequality and gives condition [1]:

$$u_s \leq u_t + \tfrac{1}{\beta^{t-1}}[\pi_t(q_s - q_t) + \rho_t^S(S_s - S_t)] \text{ for } \forall s, t = 2, ..., T$$

where $\pi_t, \rho_t^S \in \mathbb{R}_{++}$.

(2) $\implies$ (1) [*Sufficiency*] By the definition of h-rationalisable, we only need to find support prices $\pi_t \in \mathbb{R}_{++}$ that satisfy the first-order conditions given by the HAD model.

Suppose there exist numbers $u_t, S_t > 0$,$\lambda > 0$, $\rho_t^q > 0$, $\rho_t^S \leq 0$, δ, $\beta \in]0,1]$ and $\pi_t > 0$ such that

$$u_s \leq u_t + \tfrac{1}{\beta^{t-1}}[\pi_t(q_s - q_t) + \rho_t^S(S_s - S_t)] \text{ for } \forall s,t = 2,...,T \qquad [1]$$

$$\pi_t = \rho_t^q \text{ if } q_t > 0,\ t = 2,...,T \qquad [2]$$

$$\rho_t^S = (1-\delta)[\rho_t^q - \lambda\rho_t] - [\rho_{t-1}^q - \lambda\rho_{t-1}] \text{ if } q_t^k > 0,\ t = 2,...,T \qquad [3]$$

$$S_t = (1-\delta)S_{t-1} + q_{t-1} \text{ for } t = 2,...,T \qquad [4]$$

Define $\tilde{\boldsymbol{\mu}}_t = \frac{1}{\beta^{t-1}}[\pi_t, \rho_t^S]'$ and $\tilde{\boldsymbol{\chi}}_t = [q_t, S_t]'$, so the inequality can be rewritten as

$$u_s \leq u_t + \tilde{\boldsymbol{\mu}}_t'(\tilde{\boldsymbol{\chi}}_s - \tilde{\boldsymbol{\chi}}_t) \text{ for } \forall s,t = 2,....,T$$

which is a special case of the standard Afriat inequalities (where $\lambda_t = 1$ for $\forall t$). By Afriat's Theorem, the dataset $\mathbf{D}''' = \{\tilde{\boldsymbol{\mu}}_t, \tilde{\boldsymbol{\chi}}_t\}_{t=2,...,T}$ satisfies GARP and therefore there exists a locally nonsatiated, concave, continuous utility function $u(.)$ that is monotonic in $\tilde{\boldsymbol{\chi}}_t$ such that $\nabla u(\tilde{\boldsymbol{\chi}}_t) \leq \tilde{\boldsymbol{\mu}}_t$, with equality when $q_t > 0$ or $q_t = 0$ and $t \in \sigma$.

<u>Either:</u> 1. Suppose that if $q_t = 0$, then $t \in \sigma$ i.e. all observed zero purchases are due to rationing, so that $\nabla u(\tilde{\boldsymbol{\chi}}_t) = \tilde{\boldsymbol{\mu}}_t$ for $t = 2,...,T$.

Recall that $\theta_t = \pi_t - \lambda\rho_t$ if $q_t > 0$. Given π_t, λ, ρ_t^S, and $\{\theta_t\}_{t\notin\sigma}$, define $\theta_t = \sum_{i=1}^{j-1} \frac{1}{(1-\delta)^i}\rho_{t+1-i}^S + \frac{1}{(1-\delta)^j}(\rho_{t+1-j}^S + \theta_{t-j})$ if $q_t = 0$ but if $q_{t-j} > 0$, where j is the number of periods since the good was last purchased. Substituting for θ_t into condition [3] and using the definition of ρ_t^S gives the first-order condition for the stock variable:

$$\rho_t^S = \beta^{t-1}\frac{\partial u(q_t, S_t)}{\partial S_t} = (1-\delta)\theta_t - \theta_{t-1} \ ;\ t = 2,....,T$$

Given π_t, λ, and θ_t, define $\rho_t = \frac{1}{\lambda}(\pi_t - \theta_t - \gamma_t)$ if $q_t = 0$, where $0 < \gamma_t < \pi_t - \theta_t$, so that $\rho_t = \begin{cases} \frac{1}{\lambda}(\pi_t - \theta_t) & q_t > 0 \\ \frac{1}{\lambda}(\pi_t - \theta_t - \gamma_t) & q_t = 0 \end{cases}$ and $\rho_t \in \mathbb{R}_{++}$. Substituting for π_t obtains the first-order conditions for the HAD model:

$$\frac{\partial u(q_t,S_t)}{\partial q_t} = \frac{1}{\beta^{t-1}}[\lambda\rho_t + \theta_t] \text{ if } q_t > 0 \text{ ; } t = 2,..,T$$

$$\frac{\partial u(q_t,S_t)}{\partial q_t} = \frac{1}{\beta^{t-1}}[\lambda\rho_t + \theta_t + \gamma_t] \text{ if } q_t = 0 \text{ and } t \in \sigma; \ t = 2,....,T$$

Or: 2. Suppose that if $q_t = 0$, then $t \notin \sigma$ i.e. all observed zero purchases are due to corner solutions, so that $\nabla u(\tilde{\boldsymbol{\chi}}_t) \leq \tilde{\boldsymbol{\mu}}_t$, with equality when $q_t > 0$.

Recall that $\theta_t = \pi_t - \lambda\rho_t$ if $q_t > 0$. Given π_t, λ, ρ_t^S, and $\{\theta_t\}_{t\notin\sigma}$, define $\theta_t = \sum_{i=1}^{j-1} \frac{1}{(1-\delta)^i}\rho_{t+1-i}^S + \frac{1}{(1-\delta)^j}(\rho_{t+1-j}^S + \theta_{t-j})$ if $q_t = 0$ but if $q_{t-j} > 0$, where j is the number of periods since the good was last purchased. Substituting for θ_t into condition [3] and using the definition of ρ_t^S gives the first-order condition for the stock variable:

$$\rho_t^S = \beta^{t-1}\frac{\partial u(q_t,S_t)}{\partial S_t} = (1-\delta)\theta_t - \theta_{t-1} \text{ ; } t = 2,....,T$$

Given π_t, λ, and θ_t, define $\rho_t = \frac{1}{\lambda}(\pi_t - \theta_t)$ if $q_t = 0$, so $\rho_t = \frac{1}{\lambda}(\pi_t - \theta_t)$ for $\forall q_t$. Substituting for π_t obtains the first-order conditions for the HAD model:

$$\frac{\partial u(q_t,S_t)}{\partial q_t} = \frac{1}{\beta^{t-1}}[\lambda\rho_t + \theta_t] \text{ if } q_t > 0 \text{ ; } t = 2,..,T$$

$$\frac{\partial u(q_t,S_t)}{\partial q_t} \leq \frac{1}{\beta^{t-1}}[\lambda\rho_t + \theta_t] \text{ if } q_t = 0 \text{ and } t \notin \sigma; \ t = 2,....,T \ \square$$

The revealed preference conditions are a generalisation of those under fully observed prices, given in Demunyck & Verriest (2013). As in previous sections, the data can be rationalised by the HAD model even if the researcher does not observe the reason for the missing prices; in fact, observing the reason would not impose additional restrictions on the data so these conditions are again exhaustive. The support prices for zero demands are either reservation or virtual prices such that the observed choices (both non-zero and zero demands) are rational. Note that the researcher conveniently does not need to observe the stock variable in order to conduct the revealed preference test, which is often the case with consumer data.

Even though θ_t (the marginal increase in utility from a marginal increase in S_{t+1}) is unobserved by the researcher, it is not a completely free parameter but instead is determined by other parameters in the model; hence θ_t can be constructed from the solution to the linear program. If the good was purchased in period t (so ρ_t is observable), θ_t is given by the first-order condition for purchased goods and so $\theta_t = \pi_t - \lambda\rho_t$. If the good was not purchased in period t but was purchased in period $t-j$, where $1 \leq j \leq t-1$, then θ_t can be constructed using $\{\rho_t^S\}_{t=1,....,T}$ and known values of θ_t. The first-order condition for the stock variable specifies the relationship between θ_t , ρ_t^S, and θ_{t-1} that can be solved recursively to express θ_t in terms of $\{\rho_{t-i}^S\}_{i=0,...,j-1}$ and θ_{t-j} for any $j \geq 1$[21]:

$$\theta_t = \tfrac{1}{(1-\delta)}(\rho_t^S + \theta_{t-1}) = \tfrac{1}{(1-\delta)}(\rho_t^S + \tfrac{1}{(1-\delta)}(\rho_{t-1}^S + \theta_{t-2})) = ...$$

$$= \textstyle\sum_{i=1}^{j-1} \tfrac{1}{(1-\delta)^i}\rho_{t+1-i}^S + \tfrac{1}{(1-\delta)^j}(\rho_{t-j+1}^S + \theta_{t-j}) \text{ for any } 1 \leq j \leq t-1 \text{ such that } q_t > 0.$$

Since there are no intertemporal links for γ_t, γ_t cannot be constructed in a similar way. However, given the values of π_t and the constructed value of θ_t, γ_t must be within the interval $(0, \pi_t - \theta_t)$. Therefore for a given solution to the inequalities, it is possible for a set of values of γ_t to rationalise the data.

These inequalities are non-linear in unknowns, but are linear conditional on δ, β, and the initial value of the stock variable, $\bar{S}_1$. δ and β can be found by conducting a grid search in the interval $(0, 1]$. Given δ and $\bar{S}_1$, we can construct $S_2, ..., S_T$. However, $\bar{S}_1$ is usually not observed by the researcher, and cannot be constructed from the data without observing the consumption pattern of the good in all previous periods. Demuynck and Verriest (2013) suggest estimating $\bar{S}_1$ as

$$\bar{S}_1 = \tfrac{q_1}{\delta+g}$$

where $g = \frac{q_2-q_1}{q_1}$ is the estimated initial growth rate of the stock variable. This method requires the assumption that the growth rate of the stock variable in the first period observed is equal to the growth rate of actual consumption in the first period observed. As suggested

[21] The researcher can set j to be the most recent time period in which the good was purchased. Setting $j = t-1$ gives the full expression of θ_t when solved recursively until $t = 1$:

$$\theta_t = \tfrac{1}{(1-\delta)^{t-1}}(\rho_2^S + \theta_1) + \textstyle\sum_{i=0}^{t-3} \tfrac{1}{(1-\delta)^i}\rho_{t-i}^S, \text{ where } \theta_1 = \pi_1 - \lambda\rho_1$$

If $q_1 = 0$ then θ_t cannot be constructed using this expression, hence the suggestion to use later time periods in which purchases are observed.

by Demuynck and Verriest (2013), it is possible to conduct a grid search using values close to g i.e. in the interval $[g-\varepsilon, g+\varepsilon]$, where $\varepsilon < \delta$ in order for $\bar{S}_1 > 0$, which can be done as a robustness check.

The results from this section extend to the case of K storable goods i.e. $u(\mathbf{q}_t, \mathbf{S}_t)$. Each good has their own separate stock variable (S_t^k) and rate of depreciation, (δ^k). It is also possible to include both storable and non-storable goods as in Demuynck & Verriest (2013) i.e. $u(\bar{\mathbf{Q}}_t, \mathbf{q}_t, \mathbf{S}_t)$, where $\bar{\mathbf{Q}}_t$ is a vector of non-storable goods[22]; the revealed preference conditions for non-storable goods were derived in Section 1.2 and are simply added to the linear programming problem[23].

The generalisation of Theorem 1.4 to the K-good case is given below:

Theorem 1.5 (K-good HAD model). The following statements are equivalent:

(1) The data $\mathbf{D} = \{i_t, \boldsymbol{p}_t, \mathbf{q}_t\}_{t=1,\dots,T}$ is h-rationalisable.

(2) There exist numbers $u_t > 0$,$\lambda > 0$, $\{\delta^k\}_{k=1,\dots,K} \in (0,1]$, $\beta \in (0,1]$ and vectors $\mathbf{S}_t \in \mathbb{R}_+^K$, $\boldsymbol{\rho}_t^S \in \mathbb{R}_+^K$, $\boldsymbol{\rho}_t^q \in \mathbb{R}_{++}^K$, $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^K$ such that

$$u_s \leq u_t + \frac{1}{\beta^{t-1}}[\boldsymbol{\pi}_t'(\mathbf{q}_s - \mathbf{q}_t) + \boldsymbol{\rho}_t^{S'}(\mathbf{S}_s - \mathbf{S}_t)] \text{ for } \forall s,t = 2, ..., T \qquad [1]$$

$$\pi_t^k = \rho_t^{q,k} \text{ if } q_t^k > 0,\ t = 2, ..., T, k = 1, ..., K \qquad [2]$$

$$\rho_t^{S,k} = (1-\delta)[\rho_t^{q,k} - \lambda\rho_t^k] - [\rho_{t-1}^{q,k} - \lambda\rho_{t-1}^k] \text{ if } q_t^k > 0,\ t = 2, ..., T, k = 1, ..., K \qquad [3]$$

$$S_t^k = (1-\delta^k)S_{t-1}^k + q_{t-1}^k \text{ for } t = 2, ..., T, k = 1, ..., K \qquad [4]$$

[22] Although in reality all goods are storable to some extent, some goods may not be storable over the time intervals given in the data (for example, purchases of fruit in a yearly household expenditure survey).

[23] For example, support prices for both storable goods ($\boldsymbol{\pi}_t^{\boldsymbol{q}'}$) and non-storable goods ($\boldsymbol{\pi}_t^{\bar{\boldsymbol{Q}}'}$) must be found, so condition [1] from Theorem 1.5 becomes

$$u_s \leq u_t + \frac{1}{\beta^{t-1}}[\lambda\boldsymbol{\pi}_t^{\bar{\boldsymbol{Q}}'}(\mathbf{Q}_s - \mathbf{Q}_t) + \boldsymbol{\pi}_t^{\boldsymbol{q}'}(\mathbf{q}_s - \mathbf{q}_t) + \boldsymbol{\rho}_t^{S'}(\mathbf{S}_s - \mathbf{S}_t)] \text{ for } \forall s,t = 2, ..., T$$

Conditions [2]-[4] are unchanged, but there is an additional condition on the support prices for non-storable goods (given in Section 1.2):

$$\pi_t^{Q,k} = \rho_t^{Q,k} \text{ for any } Q_t^K > 0 \text{ for } \forall k = 1, ..., K, t = 1, ..., T \qquad [5]$$

Since including non-storable goods does not significantly change the revealed preference conditions or the proofs, this section will only consider the case where all goods are storable.

Proof: See Appendix 1A.

As in the 1-good case, the non-linear inequalities can be made linear by conditioning on $\boldsymbol{\delta}$, β, and $\bar{\mathbf{S}}_1$, where $\bar{S}_1^k$ is defined as before for $\forall k$. However, conducting a grid search over $K+1$ dimensions for the values $\{\beta, \delta^1, ..., \delta^k\}$ would be extremely computationally intensive for large K, so the HAD model may not be recommended for these cases.

1.3.3 Links between the models

The life cycle model, SMH model, and HAD model vary in the restrictions imposed on the data, and some forms of these models are nested. The differences between these models do not depend on whether prices are fully observed or not. The life cycle model is the most restrictive due to the consumption independence assumption, so if a dataset is consistent with the life cycle model, it is consistent with the other two models. The life cycle model is a special case of the HAD model where the stock variable S_t is absent from the utility function. The life cycle model is also nested in the R-lag SMH model, which can be seen by setting the discounted shadow prices $\tilde{\rho}_t^{t+r,k}$ as 0 for $\forall r \geq 1$ i.e. no intertemporal links due to consumption of the good[24]. Since the life cycle model implies more restrictions than GARP (expenditure efficiency both within and across periods rather than only within periods), if the data is rationalisable by the R-lag SMH model with the additional restriction that $\tilde{\rho}_t^{t+r,k} = 0$ $(r \geq 1)$, then the data also satisfies GARP[25]. However, without this additional restriction, data that is rationalisable by the R-lag SMH model need not be rationalisable by GARP.

In general, the R-lag SMH model is independent from the HAD model, with the exception of $R = 1$, which is a special case of the HAD model where $\delta = 1$ (complete depreciation of last period's stock) so that $S_t = q_{t-1}$ (Demunyck & Verriest, 2013). Another difference is that there are $T-1$ distinct inequalities for the HAD model compared to $T-R$ for the

[24]See Theorem 2 in Crawford (2010) for a formal proof of this result.
[25]See Corollary 1 in Crawford (2010).

R-lag SMH model, which may matter for empirical applications, depending on the size of T. Demuynck and Verriest (2013) explain that neither model is more restrictive than the other: the SMH-model restricts the effects of current consumption to a finite number of future periods instead of all future periods, but does not specify how current consumption affects future welfare, whereas the HAD model allows current consumption to affect all future periods but restricts the intertemporal utility effects to the stock variable S_t. Therefore, in general the R-lag SMH model and HAD model are non-nested.

Storability of goods can be modelled in other ways besides habits models. Instead of having past consumption enter into the utility function (either directly or indirectly via a stock variable) and drawing an analogy between storability and habits, buying storable goods can be likened to asset investment, as in the model of Deaton (1992). Unlike the investment model of Demuynck and Verriest (2013), goods ('assets') appear in the budget constraint but are not included in the utility function, so the consumer does not get utility directly from the stock of storable goods, but can get utility indirectly by selling some of his stock to buy goods in the current period. Appendix 1B shows that this alternative way of modelling storability also gives rise to revealed preference conditions based only on observables that can be tested by solving a linear programming problem, so the general conclusions from Section 1.3 are robust to the way in which storability is modelled.

Additional remarks about the results in Sections 1.2 and 1.3

1. For data with large amount of goods, results only hold for the subset of goods that are purchased at least once. Varian (1988) shows that if total expenditure is unobserved, either because a good is not purchased at all (we observe the quantity purchased but not its price) or because we observe the price but not the quantity purchased, then revealed preference

theory puts no restrictions on the data unless we make separability assumptions[26]. Hence all revealed preference tests, including those in this chapter, are joint tests of separability and utility maximisation, regardless of whether some prices are missing or not[27].

2. While it is obvious that the results in this chapter are valid for the extreme case in which all prices are observed (in which case the standard theorems apply), they also hold for the other extreme case in which none of the prices are observed. In this case, rationalisability depends on finding support prices for all goods that are purchased at least once.

3. As with the tests for rationality, the set of support prices constructed from the solution to the revealed preference test also does not depend on the reason for missing prices. For example, in the life cycle model, if good k is not purchased, then the constructed discounted price $\rho_t^k \in (0, \pi_t^k)$, either under the assumption that the zero purchase was due to rationing or a corner solution.

4. For all models used in this chapter, the marginal disutility of rationing γ_t^k is set identified and depends on the parameters from the revealed preference test[28].

1.4 Examples

The following examples demonstrate that datasets are still falsifiable and rationalisable when prices are partially observed; that is, it is possible to find data which are rationalisable by each of the models in Sections 1.2 and 1.3, and data which are not.

[26]Theorem 1 in Varian (1988) states that if we observe the quantity purchased of a particular good but not its price, then there always exists a set of prices for that good such that the data satisfy GARP. Similarly, Theorem 2 in Varian (1988) states that if we observe the prices of the omitted good but not the quantities purchased, then there always exists a set of quantities such that the data satisfy GARP.

[27]It is possible to relax this separability assumption: Deb et al (2017) develop revealed preference conditions for the case where total expenditure is not observed, without needing to make the weak separability assumption implicit in Afriat's Theorem.

[28]For the life cycle model, $\gamma_t^k \in (0, \lambda\pi_t^k)$; for the SMH R-lag model, $\gamma_t^k \in (0, \lambda[\beta^{t-1}\tilde{\pi}_t^k + \sum_{i=1}^{R} \rho_t^{t+i,k})$; for the HAD model, $\gamma_t^k \in (0, \pi_t^k - \theta_t^k)$.

1.4.1 Life cycle model

Example 1: $(T = 2, K = 3)$

Pass: $\mathbf{p} = \begin{bmatrix} 2 & 4 \\ 3 & 4 \\ 2 & . \end{bmatrix}, \mathbf{q} = \begin{bmatrix} 3 & 1 \\ 5 & 1 \\ 3 & 0 \end{bmatrix}$ for $\beta \in [\frac{12}{11(1+i)}, 1]$

Fail: $\mathbf{p} = \begin{bmatrix} 2 & 4 \\ 3 & 4 \\ . & 3 \end{bmatrix}, \mathbf{q} = \begin{bmatrix} 1 & 3 \\ 1 & 5 \\ 0 & 3 \end{bmatrix}$

For the simple case above, it is easy to show why these datasets pass or fail using direct calculation:

From Theorem 1.1, rationalisability by the life cycle model requires that $u_s \leq u_t + \frac{1}{\beta^{t-1}}[\lambda\boldsymbol{\pi}_t'(\mathbf{q}_s - \mathbf{q}_t)]$ for $\forall s, t = 1, 2$, where $\pi_t^k = \rho_t^k$ for any $q_t^k > 0$ for $k = 1, 2, 3; t = 1, 2$. Summing over all s, t gives the requirement $0 \leq \sum_{s,t\in\{1,2\}} \frac{1}{\beta^{t-1}}\boldsymbol{\pi}_t(\mathbf{q}_s - \mathbf{q}_t)$. The data thus violates the life cycle model if we cannot find $\beta \in (0, 1]$ and an interest rate i such that this inequality holds. Substituting in the $\mathbf{p}$ and $\mathbf{q}$ given in the first case gives $\frac{1}{\beta(1+i)}[4(2) + 4(4)] + [2(-2) + 3(-4) + 2(-3)]$, which satisfies the requirement if $\beta \geq \frac{12}{11(1+i)}$. Since $\beta \in (0, 1]$ and $i \ll 1$ in plausible situations ($i = 0.05$ is a 5% interest rate between periods), it is possible to find β, i such that the requirement holds. One example is $i = 10\%$, so that $\beta = 0.99$, which is reasonable for annual data. However, substituting in the $\mathbf{p}$ and $\mathbf{q}$ given in the second case gives $\frac{1}{\beta(1+i)}[4(-2) + 4(-4)] + [2(2) + 3(4)]$, which satisfies the requirement if $\beta \geq \frac{3}{2(1+i)}$, where i is the interest rate in period 2. For $\beta \in (0, 1]$, the interest rate must be larger than 1.5 (150% per period), which is implausible even if the data is annual. Since there is no $\beta \in (0, 1]$ that satisfies the revealed preference conditions, by Theorem 1.1, the data is not rationalisable by the life cycle model for plausible interest rates in the second case[29].

[29]It is possible to rationalise the data for plausible interest rates by allowing $\beta > 1$, for example if $i = 0.1$, then $\beta = 1.36$, or if $i = 0.05$, then $\beta = 1.43$. It is then up to the researcher to decide

The data from Example 1 that failed the life cycle model test passes the static model test of Crawford and Polisson (2015) which allows for the marginal utility of income λ to vary across periods (the consumer maximises according to a per-period budget constraint rather than a lifetime budget constraint). The revealed preference conditions of Crawford and Polisson imply that $0 \leq \sum_{s,t\in\{1,2\}} \boldsymbol{\pi}_t(\mathbf{q}_s - \mathbf{q}_t)$, where $\pi_t^k = \lambda_t p_t^k$ for any $q_t^k > 0$ for $k = 1,2,3; t = 1,2$. Substituting in for $\mathbf{p}$ and $\mathbf{q}$ gives $\lambda_2[4(-3) + 4(-4)] + \lambda_1[2(3) + 3(4)]$, and it is possible to choose λ_1, λ_2 such that this statement is greater than 0. For example, setting $\lambda_1 = 1.5, \lambda_2 = 0.5$ (consumer's per-period income at $t = 1$ is lower than at $t = 2$) can explain why the consumer purchases more of each good at $t = 2$ even though the prices of all goods are higher.

The intuition behind this example is as follows: according to the life cycle model with perfect foresight, the consumer should have accounted for the higher prices at $t = 2$ (which are still higher after discounting, for plausible interest rates e.g. 5%) when deciding how much to purchase at $t = 1$, so it is not possible to rationalise purchasing more when prices are high and purchasing almost nothing when prices are low. However, according to the static model, the consumer only accounts for his income in that period, so it is possible to rationalise this behaviour by attributing it to differences in the consumer's income in each period (if the consumer has a much larger budget at $t = 2$ than at $t = 1$ then it makes sense to purchase more even if prices are higher).

Rationalisability in dynamic models depends on both the timing of purchases (due to discounting on future prices and the weight assigned to future utility) and the price at which goods were purchased, whereas in static models it only depends on the latter. Due to consumption independence across periods, rationalisability depends more on prices than timing: the data in Example 1 that failed the life cycle test can be made to pass by switching the

whether or not these values of β are reasonable representations of consumer behaviour.

columns of $\mathbf{q}$ only (purchasing more when prices are low, and less when prices are high)[30], but not by switching both columns of $\mathbf{p}$ and $\mathbf{q}$ (purchasing the same price-quantity combinations in different time periods).

1.4.2 Storable goods models

SMH model

Consider $R = 1$ (purchases only last for 2 periods) and a 5% interest rate.

Example 2: $(T = 4, K = 2)$

$$\text{Pass: } \mathbf{p} = \begin{bmatrix} 5 & . & . & . \\ . & 1 & . & 5 \end{bmatrix}, \mathbf{q} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 4 & 0 & 1 \end{bmatrix} \text{ for } \beta = 0.01^{31}$$

$$\text{Fail: } \mathbf{p} = \begin{bmatrix} 5 & . & . & . \\ . & 1 & 5 & . \end{bmatrix}, \mathbf{q} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 4 & 1 & 0 \end{bmatrix}$$

Example 3: $(T = 4, K = 2)$

$$\text{Pass: } \mathbf{p} = \begin{bmatrix} 5 & . & . & . \\ . & . & 5 & 1 \end{bmatrix}, \mathbf{q} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 4 \end{bmatrix}$$

$$\text{Fail: } \mathbf{p} = \begin{bmatrix} 5 & . & . & . \\ . & . & 5 & 3 \end{bmatrix}, \mathbf{q} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 4 \end{bmatrix} \text{ for } \beta = 0.01$$

As with the life cycle model, both the timing of purchases and the relative price at which goods were purchased are important. However, due to consumption dependence, the revealed preference test may be more sensitive to the timing of purchases than the purchase price, because only certain patterns in the zero demands are consistent with storability (as defined in the model). In Example 2, buying the same amount of good 2 in a different time period at the same discounted price violates the model, while in Example 3, buying the

[30] The intuition behind the lower bound on β is that β is a measure of how much the consumer values future consumption relative to current consumption and so must be sufficiently high in order for the consumer to leave some of his lifetime income for future spending.

[31] The grid search used started with $\beta = 0.01$, with a spacing of 0.01, so the lowest value of β that rationalises the data is listed here.

same amount of good 2 in the same time periods but at a higher price violates the model. These examples show that even if the data contains many missing observations, the SMH model still imposes enough restrictions on support prices so that it is possible to find data that violates the model.

In both examples, the data failed when only one restriction on support prices was imposed (from Theorem 1.2, the condition on discounted support prices only needs to hold for nonzero demands in adjacent periods). Since each good lasts for two periods, the second dataset in Example 2 fails because the consumer bought good 2 when he still had some (purchased in the previous period) stored at home. Given the prices in the example, it is not possible to find a discounted willingness to pay for current and future consumption that satisfies both the intertemporal restrictions and the Afriat-type inequalities. While both datasets in Example 3 have one restriction on support prices, the second dataset in Example 3 fails because the price at $t = 4$ was too high to satisfy the Afriat-type inequalities.

HAD model

To focus on the intuition, consider $K = 1$ and a 5% interest rate. Since the revealed preference conditions for one good ([2]-[4] in Theorems 1.4 and 1.5) do not depend on other goods, data will fail the test if there is a violation in only one good.

Example 4:

Pass[32] : $\mathbf{p} = \begin{bmatrix} 3 & 3 & . & . \end{bmatrix}, \mathbf{q} = \begin{bmatrix} 3 & 1 & 0 & 0 \end{bmatrix}$ for $\beta = 1, \delta = 0.95$

Fail: $\mathbf{p} = \begin{bmatrix} . & . & 3 & 3 \end{bmatrix}, \mathbf{q} = \begin{bmatrix} 0 & 0 & 3 & 1 \end{bmatrix}$

Example 5:

Pass: $\mathbf{p} = \begin{bmatrix} 10 & 3 & . & . \end{bmatrix}, \mathbf{q} = \begin{bmatrix} 5 & 1 & 0 & 0 \end{bmatrix}$ for $\beta = 1, \delta = 0.95$

[32] There are other combinations of β and δ that work, but since the range of δ varies according to the value of β, it would be cumbersome to list all combinations here. Instead, the highest value of β and the corresponding highest value of δ that rationalises the data will be listed (the grid search used starts with $\beta = 1$ and checks all values of δ starting with $\delta = 1$, with a spacing of 0.05).

Fail: $\mathbf{p} = \begin{bmatrix} . & . & 10 & 3 \end{bmatrix}, \mathbf{q} = \begin{bmatrix} 0 & 0 & 5 & 1 \end{bmatrix}$

Example 6 (consumption in non-consecutive periods):

Pass: $\mathbf{p} = \begin{bmatrix} . & 3 & . & 10 \end{bmatrix}, \mathbf{q} = \begin{bmatrix} 0 & 1 & 0 & 5 \end{bmatrix}$ for $\beta = 1, \delta = 1$

Pass: $\mathbf{p} = \begin{bmatrix} 10 & . & 3 & . \end{bmatrix}, \mathbf{q} = \begin{bmatrix} 5 & 0 & 1 & 0 \end{bmatrix}$ for $\beta = 1, \delta = 1$

As with the SMH model, these three examples show that both the timing of purchases and the price at which goods were purchased are important, but unlike the life cycle model, rationalisability depends more on timing because of consumption dependence via the unobserved stock variable. Compared to the SMH model, in which purchases of goods only last for $R+1$ periods, in the HAD model purchases in all earlier periods contribute to the (unobserved) stock variable at any given time period observed in the data (the degree of contribution depends on the depreciation rate δ). As discussed in Section 1.3.3, the SMH and HAD model are non-nested in general, so it is possible for some consumption patterns to be consistent with storability according to one habits model but not according to the other. In fact, all of the above datasets pass the 1-lag SMH model for some plausible value of β.

In Example 4, the given purchasing pattern (prices and quantities) passes the HAD model if done in periods 1 and 2, but fails the HAD model if done in periods 3 and 4. Part of the reason for this is that it is not possible to construct a stock variable that rationalises the data. In particular, the stock in all future periods depends on δ and $\bar{S}_1$ (a function of q_1 and q_2). With $q_1, q_2 = 0$, the initial stock at $t = 1$ ($\bar{S}_1$) is estimated to be 0, and $S_2 = 0$ also. Therefore it is not possible to rationalise non-purchase in the first two periods by saying that the consumer already has some of that good stored at home. With $q_1, q_2 \neq 0$, the consumer has built up a stock of the good (nonzero initial stock and next period's stock), which can explain the zero demand in future periods, even if the stock depreciates at a high rate. The datasets in Example 5 pass/fail the HAD model for the same reasons as those in Example 4, but also demonstrate that it is possible to rationalise purchases at high prices (purchasing when the price is high and not purchasing when the price is low can violate rationality, but does not always have to). Here, the high value assigned to future consumption (β) means it can be beneficial to purchase a lot in earlier periods, even at high prices, in order to build up the stock of the good and receive utility from it in future periods. Since current consumption affects the stock variable in all future periods, the large amounts of initial purchases can explain why the consumer does not purchase anything in later period, even with a high depreciation rate.

Example 6 shows that as with the 1-lag SMH model, restrictions are only imposed on the data if consumption is observed in consecutive periods. The reason for this condition is more obvious for the 1-lag SMH model, because the condition on support prices only holds if purchases are observed in adjacent periods. For the HAD model, this condition is due to the additional parameter δ, which determines what proportion of the current stock is available next period. As discussed in Section 1.3.3, setting $\delta = 1$ makes the HAD model equivalent to the 1-lag SMH model, so any purchases after a gap can be explained by having the stock variable depreciate completely. In Example 6, both datasets pass regardless of the price/quantity combinations used, because by the time the consumer purchases the good again, they have none left at home ($S_t = q_{t-1}$).

1.4.3 Discussion

Table 1.1 summarises some general characteristics of data that pass or fail each of the models considered[33], based on the examples in Sections 1.4.1 and 1.4.2:

[33]Note that when using actual data, the test result depends on the particular data values and the interest rate.

Table 1.1: Rationalisability: General characteristics of data

Model	Pass	Fail
Life cycle	Buying smaller amounts when price is high and larger amounts when price is low.	Buying larger amounts when price is high and smaller amounts when price is low.
SMH	Gaps between purchases larger than the lag length (subsequent purchases can be rationalised because the consumer's stock has run out).	For purchases in consecutive periods, buying large amounts when price is high and small amounts when price is low.
HAD	Gaps (of any length) between purchases ($\delta = 1$ can be chosen so that stock variable depreciates completely during the gap in purchasing); Purchasing large amounts then not purchasing anything later.	Not purchasing anything in the first few periods and then purchasing later (difficult to construct a stock variable that rationalises such behaviour).

Some of the models can be used to make demand predictions and conduct welfare analysis, even with partially observed prices. Crawford and Polisson (2015) show that for static utility maximisation, the convexity of the set of support prices and the support set is preserved under partially observed prices[34]. This result applies to the life cycle model, which is a special case of the static utility maximisation model, but does not always hold for the storable goods models[35]. Therefore, researchers may prefer to use the life cycle model when there are missing prices.

There are some natural extensions to the work done in this chapter. The life cycle results could be used for demand predictions and welfare analysis in an empirical setting, for example using data from developing countries, in which price information is not always available. While the examples in Sections 1.4.1 and 1.4.2 demonstrate that the models are still falsifiable when prices are partially observed, there are fewer restrictions on behaviour so

[34]See Proposition 2 in Crawford and Polisson (2015) for further discussion and proof.

[35]The SMH model is not convex, as the following counterexample demonstrates:

$$\mathbf{q} = \begin{bmatrix} 4 & 1 & 5 & 2 \\ 4 & 4 & 2 & 4 \end{bmatrix}, \mathbf{p}_1 = \begin{bmatrix} 5 & 4 & 5 & 5 \\ 5 & 4 & 2 & 5 \end{bmatrix}, \mathbf{p}_2 = \begin{bmatrix} 1 & 3 & 2 & 3 \\ 5 & 4 & 2 & 5 \end{bmatrix}, i = 5\%$$

The convexity of the HAD model is still an open question, due to the difficulty of either proving convexity or finding a counterexample.

it is likely that power (the probability of finding data that fails the revealed preference test) would decrease. The degree to which partially observed prices decreases the power and widens the bounds on demand predictions and welfare measures would be worth investigating further.

1.5 Conclusion

This chapter has derived necessary and sufficient conditions for data with missing price observations to be consistent with dynamic models, allowing for goods to be storable. Zero purchases in some periods can be due to corner solutions, rationing, or the consumer could already have a stored stock of the good. However, the nonparametric revealed preference tests suggested for the models in Sections 1.2 and 1.3 do not require the researcher to observe which of these three reasons is responsible for the missing prices or observe additional variables aside from those already given in the dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,\ldots,T}$. It is also not necessary to impute prices in order to conduct the revealed preference tests; in fact, if a solution to the given inequalities exists, actual prices can be constructed from the support prices according to the particular reason. Any unobservable parameters in the models considered (for example, the marginal disutility from rationing γ_t^k in all models, or the marginal utility from next period's stock θ_t in the HAD model) can be also constructed from the solution. The proofs of Theorems 1-5 showed that these revealed preference conditions are also exhaustive, so even knowing which reason(s) resulted in the missing prices would not impose additional restrictions on the data. Section 1.4 presents examples of datasets with missing prices that are not consistent with these models, demonstrating that even if actual prices are partially observed, the revealed preference conditions are still falsifiable due to the conditions given for the support prices and discounted shadow prices.

Appendix 1A: Proofs

Proof of Theorem 1.3 (R-lag SMH model):

(1) $\implies$ (2) [*Necessity*] The dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ is sr-rationalisable if there exists a locally non-satiated, continuous, concave, and differentiable utility function $u(.)$ and numbers $\lambda > 0$, $\gamma_t^k > 0$, $\beta \in (0, 1]$ such that

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1},...,\mathbf{q}_{t-R})}{\partial q_t^k} + \sum_{r=1}^{R} \beta^{t+r-1}\frac{\partial u(\mathbf{q}_{t+r},\mathbf{q}_{t+r-1},...,\mathbf{q}_{t+r-R})}{\partial q_t^k} = \lambda\rho_t^k \text{ if } q_t^k > 0$$

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1},...,\mathbf{q}_{t-R})}{\partial q_t^k} + \sum_{r=1}^{R} \beta^{t+r-1}\frac{\partial u(\mathbf{q}_{t+r},\mathbf{q}_{t+r-1},...,\mathbf{q}_{t+r-R})}{\partial q_t^k} \leq \lambda\rho_t^k \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t$$

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1},...,\mathbf{q}_{t-R})}{\partial q_t^k} + \sum_{r=1}^{R} \beta^{t+r-1}\frac{\partial u(\mathbf{q}_{t+r},\mathbf{q}_{t+r-1},...,\mathbf{q}_{t+r-R})}{\partial q_t^k} = \lambda\rho_t^k + \gamma_t^k \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t$$

These first-order conditions can be rearranged in terms of $\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1},...,\mathbf{q}_{t-R})}{\partial q_t^k}$:

$$\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1},...,\mathbf{q}_{t-R})}{\partial q_t^k} = \frac{1}{\beta^{t-1}}\lambda\tilde{\rho}_t^{t,k} \text{ if } q_t^k > 0 \text{ ; } t = R+1,...,T, k = 1,...,K$$

$$\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1},...,\mathbf{q}_{t-R})}{\partial q_t^k} \leq \frac{1}{\beta^{t-1}}\lambda\tilde{\rho}_t^{t,k} \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t \text{ ; } t = R+1,...,T, k = 1,...,K$$

$$\frac{\partial u(\mathbf{q}_t,\mathbf{q}_{t-1},...,\mathbf{q}_{t-R})}{\partial q_t^k} = \frac{1}{\beta^{t-1}}\lambda[\tilde{\rho}_t^{t,k} + \frac{\gamma_k^k}{\lambda}] \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t; \ t = R+1,...,T, k = 1,...,K$$

where $\tilde{\rho}_t^{t,k} = \rho_t^k - \sum_{r=1}^{R} \tilde{\rho}_t^{t+r,k}$.

By concavity of $u(.)$,

$$u_s \leq u_t + \nabla u(\boldsymbol{\chi}_t)'(\boldsymbol{\chi}_s - \boldsymbol{\chi}_t) \text{ for } \forall s, t = 1, ..., T$$

where $\boldsymbol{\chi}_t^k = [q_t^k, q_{t-1}^k, ..., q_{t-R}^k]'$.

For $t = R+1, ..., T, k = 1, ..., K$, define $\tilde{\pi}_t^k = \frac{1}{\beta^{t-1}}\tilde{\rho}_t^{t,k}$ if $q_t^k > 0$ or $q_t^k = 0$ and $k \notin \kappa_t$, and $\tilde{\pi}_t^k = \frac{1}{\beta^{t-1}}(\tilde{\rho}_t^{t,k} + \frac{\gamma_t^k}{\lambda})$ if $q_t^k = 0$ and $k \in \kappa_t$. Define $\boldsymbol{\pi}_t^k = [\tilde{\pi}_t^k, \tilde{\pi}_{t-1}^k, ..., \tilde{\pi}_{t-R}^k]'$. Note that $\boldsymbol{\pi}_t^k \in \mathbb{R}_{++}^{R+1}$ since $\tilde{\rho}_t^{t,k}$, γ_t^k, and $\lambda > 0$, so $\nabla u(\boldsymbol{\chi}_t) \leq \lambda\boldsymbol{\pi}_t$, with equality if $q_t^k > 0$. Substituting $\lambda\boldsymbol{\pi}_t$ for $\nabla u(\mathbf{q}_t)$ therefore preserves the inequality and gives statement (2):

$$u_s \leq u_t + \lambda\boldsymbol{\pi}_t'(\boldsymbol{\chi}_s - \boldsymbol{\chi}_t) \text{ for } \forall s, t = R+1, ..., T$$

$$\rho_t^k = \sum_{r=0}^{R} \tilde{\rho}_t^{t+r,k} \text{for any } q_t^k > 0, \ k = 1, ..., K, \ t \in \{R+1, ..., T\}$$

$$\boldsymbol{\pi}_t^k = \frac{1}{\beta^{t-1}}[\tilde{\rho}_t^{t,k}, \tilde{\rho}_{t-1}^{t,k}, ..., \tilde{\rho}_{t-R}^{t,k}]' \text{ for any } \boldsymbol{\chi}_t^k > 0,\ k = 1, ..., K, t \in \{R+1, ..., T\}$$

where $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^{K(R+1)}$.

(2) $\Longrightarrow$ (1) [*Sufficiency*] By the definition of sr-rationalisable, we only need to find support prices $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^{K(R+1)}$ that satisfy the first-order conditions given by the R-lag SMH model.

Suppose there exist numbers $u_t \in \mathbb{R}$, $\lambda > 0$, $\beta \in (0, 1]$, $\{\tilde{\rho}_{t-i}^{t,k}\}_{i=0,...,R}^{k=1,...,K} > 0$, $\{\tilde{\rho}_t^{t+i,k}\}^{i=0,..,R;k=1,...,K} \neq 0$, and vectors $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^{K(R+1)}$ such that

$$u_s \leq u_t + \lambda \boldsymbol{\pi}_t'(\boldsymbol{\chi}_s - \boldsymbol{\chi}_t) \text{ for } \forall s, t = R+1, ..., T$$

$$\rho_t^k = \sum_{r=0}^{R} \tilde{\rho}_t^{t+r,k} \text{for any } q_t^k > 0,\ k = 1, ..., K,\ t \in \{R+1, ..., T\}$$

$$\boldsymbol{\pi}_t^k = \frac{1}{\beta^{t-1}}[\tilde{\rho}_t^{t,k}, \tilde{\rho}_{t-1}^{t,k}, ..., \tilde{\rho}_{t-R}^{t,k}]' \text{ for any } \boldsymbol{\chi}_t^k > 0,\ k = 1, ..., K, t \in \{R+1, ..., T\}$$

The inequalities $\{u_s \leq u_t + \lambda \boldsymbol{\pi}_t^{'}(\boldsymbol{\chi}_s - \boldsymbol{\chi}_t)\}_{s,t=R+1,...,T}$ are a special case of the standard Afriat inequalities (where $\lambda_t = \lambda$ for $\forall t$). By Afriat's Theorem, the dataset $\mathbf{D} = \{\boldsymbol{\pi}_t, \boldsymbol{\chi}_t\}_{t=1,...,T}$ satisfies GARP and therefore there exists a locally nonsatiated, concave, continuous utility function $u(.)$ that is monotonic in $\boldsymbol{\chi}_t$ such that $\nabla u(\boldsymbol{\chi}_t) \leq \lambda \boldsymbol{\pi}_t$, with equality when $q_t^k > 0$ or $q_t^k = 0$ and $k \in \kappa_t$.

<u>Either:</u> 1. Suppose that if $q_t^k = 0$, then $k \in \kappa_t$ i.e. all observed zero purchases are due to rationing, so that $\nabla u(\boldsymbol{\chi}_t) = \lambda \boldsymbol{\pi}_t$ for $t = R+1, ..., T$.

Consider the first row of $\nabla_k u(\boldsymbol{\chi}_t)$: $\frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1}, ..., \mathbf{q}_{t-R})}{\partial q_t^k} = \lambda \tilde{\pi}_t^k$. Given $\tilde{\pi}_t^k$, $\{\tilde{\rho}_t^{t+i,k}\}^{i=0,...,R;k=1,...,K}$, β, and λ, define $\tilde{\rho}_t^k = \beta^{t-1}\tilde{\pi}_t^k + \sum_{i=1}^{R} \tilde{\rho}_t^{t+i,k} - \frac{\gamma_t^k}{\lambda}$ if $q_t^k = 0$, where $0 < \gamma_t^k < \lambda[\beta^{t-1}\tilde{\pi}_t^k + \sum_{i=1}^{R} \tilde{\rho}_t^{t+i,k}]$, so that $\rho_t^k = \begin{cases} \beta^{t-1}\tilde{\pi}_t^k + \sum_{i=1}^{R} \tilde{\rho}_t^{t+i,k} & q_t^k > 0 \\ \beta^{t-1}\tilde{\pi}_t^k + \sum_{i=1}^{R} \tilde{\rho}_t^{t+i,k} - \frac{\gamma_t^k}{\lambda} & q_t^k = 0 \end{cases}$ and $\boldsymbol{\rho}_t \in \mathbb{R}_{++}^{K(R+1)}$. Substituting for $\tilde{\pi}_t^k$ obtains the first-order conditions for the R-lag SMH life cycle model:

$$\frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1}, ..., \mathbf{q}_{t-R})}{\partial q_t^k} = \frac{1}{\beta^{t-1}} \lambda \tilde{\rho}_t^{t,k} \text{ if } q_t^k > 0 \text{ ; } t = R+1, ..., T, k = 1, ..., K$$

$$\frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1}, ..., \mathbf{q}_{t-R})}{\partial q_t^k} = \frac{1}{\beta^{t-1}} \lambda [\tilde{\rho}_t^{t,k} + \frac{\gamma_k^k}{\lambda}] \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t;\ t = R+1, ..., T, k = 1, ..., K$$

where $\tilde{\rho}_t^{t,k} = \rho_t^k - \sum_{r=1}^R \tilde{\rho}_t^{t+r,k}$.

Or: 2. Suppose that if $q_t^k = 0$, then $k \notin \kappa_t$ i.e. all observed zero purchases are due to corner solutions, so that $\nabla u(\boldsymbol{\chi}_t) \leq \lambda \boldsymbol{\pi}_t$, with equality when $q_t^k > 0$.

Consider the first row of $\nabla_k u(\boldsymbol{\chi}_t)$: $\frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1}, \ldots, \mathbf{q}_{t-R})}{\partial q_t^k} < \lambda \tilde{\pi}_t^k$, with equality when $q_t^k > 0$.

Given $\tilde{\pi}_t^k$, define $\rho_t^k = \beta^{t-1}\tilde{\pi}_t^k + \sum_{i=1}^R \tilde{\rho}_t^{t+i,k}$ for $q_t^k = 0$, so $\rho_t^k = \beta^{t-1}\tilde{\pi}_t^k + \sum_{i=1}^R \tilde{\rho}_t^{t+i,k}$ for $\forall q_t^k$.

Substituting for $\tilde{\pi}_t^k$ obtains the first-order conditions for the R-lag SMH model:

$$\frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1}, \ldots, \mathbf{q}_{t-R})}{\partial q_t^k} = \frac{1}{\beta^{t-1}} \lambda \tilde{\rho}_t^{t,k} \text{ if } q_t^k > 0 \text{ ; } t = R+1, ..., T, k = 1, ..., K$$

$$\frac{\partial u(\mathbf{q}_t, \mathbf{q}_{t-1}, \ldots, \mathbf{q}_{t-R})}{\partial q_t^k} \leq \frac{1}{\beta^{t-1}} \lambda \tilde{\rho}_t^{t,k} \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t \text{ ; } t = R+1, ..., T, k = 1, ..., K$$

where $\tilde{\rho}_t^{t,k} = \rho_t^k - \sum_{r=1}^R \tilde{\rho}_t^{t+r,k}$.

Proof of Theorem 1.5 (HAD model with K goods):

(1) $\implies$ (2) [*Necessity*] The dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,\ldots,T}$ is h-rationalisable if there exists a locally non-satiated, continuous, concave, and differentiable utility function $u(.)$ and numbers $\lambda > 0$, $\theta_t^k > 0$, $\gamma_t^k \in \mathbb{R}$, $\{\delta^k\}_{k=1,\ldots,K} \in (0,1]$, and $\beta \in (0,1]$ such that

$$\frac{\partial u(\mathbf{q}_t, \mathbf{S}_t)}{\partial q_t^k} = \frac{1}{\beta^{t-1}}[\lambda \rho_t^k + \theta_t^k] \text{ if } q_t^k > 0 \text{ ; } t = 2, ..., T, k = 1, ..., K$$

$$\frac{\partial u(\mathbf{q}_t, \mathbf{S}_t)}{\partial q_t^k} \leq \frac{1}{\beta^{t-1}}[\lambda \rho_t^k + \theta_t^k] \text{ if } q_t^k = 0 \text{ and } k \notin K_t \text{ ; } t = 2, ..., T, k = 1, ..., K$$

$$\frac{\partial u(\mathbf{q}_t, \mathbf{S}_t)}{\partial q_t^k} = \frac{1}{\beta^{t-1}}[\lambda \rho_t^k + \theta_t^k + \gamma_t^k] \text{ if } q_t^k = 0 \text{ and } k \in K_t; \; t = 2, ..., T, k = 1, ..., K$$

$$\frac{\partial u(\mathbf{q}_t, \mathbf{S}_t)}{\partial S_t^k} = \frac{1}{\beta^{t-1}}[(1-\delta^k)\theta_t^k - \theta_{t-1}^k] \text{ ; } t = 2, ..., T, k = 1, ..., K$$

Conditions [3] is obtained by substituting for θ_t^k and θ_{t-1}^k into the first-order condition for stock variable, and using the definitions of the discounted shadow prices $\rho_t^{S,k}$ and $\rho_t^{q,k}$. Condition [4] is the equation of motion for the stock variable.

By concavity of $u(.)$,

$$u_s \leq u_t + \nabla_{\mathbf{q}_t} u(\mathbf{q}_t, \mathbf{S_t})'(\mathbf{q}_s - \mathbf{q}_t) + \nabla_{\mathbf{S}_t} u(\mathbf{q}_t, \mathbf{S}_t)'(\mathbf{S}_s - \mathbf{S}_t) \text{ for } \forall s, t = 2, ..., T$$

For $t = 2, ..., T, k = 1, ..., K$, define $\pi_t^k = \lambda\rho_t^k + \theta_t^k$ if $q_t^k > 0$ or $q_t^k = 0$ and $k \notin K_t$, and $\pi_t^k = \lambda\rho_t^k + \theta_t^k + \gamma_t^k$ if $q_t^k = 0$ and $k \in K_t$. Note that $\pi_t^k = \beta^{t-1}\frac{\partial u(\mathbf{q}_t, \mathbf{S}_t)}{\partial q_t^k} = \rho_t^{q,k}$ if $q_t^k > 0$, so condition [2] holds. $\pi_t^k > 0$ since $\theta_t^k \geq 0$, and ρ_t^k, γ_t^k, $\lambda > 0$, so $\nabla_{\mathbf{q}_t} u(\mathbf{q}_t, \mathbf{S_t}) \leq \frac{1}{\beta^{t-1}}\boldsymbol{\pi}_t$, with equality if $q_t^k > 0$. Also define $\rho_t^{S,k} = (1-\delta^k)\theta_t^k - \theta_{t-1}^k$ for $t = 2, ..., T, k = 1, ..., K$, so $\nabla_{\mathbf{S}_t} u(\mathbf{q}_t, \mathbf{S_t}) = \frac{1}{\beta^{t-1}}\boldsymbol{\rho}_t^S$.

Substituting $\frac{1}{\beta^{t-1}}\boldsymbol{\pi}_t$ for $\nabla_{\mathbf{q}_t} u(\mathbf{q}_t, \mathbf{S}_t)$ and $\frac{1}{\beta^{t-1}}\boldsymbol{\rho}_t$ for $\nabla_{\mathbf{S}_t} u(\mathbf{q}_t, \mathbf{S_t})$ therefore preserves the inequality and gives condition [1]:

$$u_s \leq u_t + \frac{1}{\beta^{t-1}}[\boldsymbol{\pi}_t'(\mathbf{q}_s - \mathbf{q}_t) + \boldsymbol{\rho}_t^{S'}(\mathbf{S}_s - \mathbf{S}_t)] \text{ for } \forall s, t = 2, ..., T$$

where $\boldsymbol{\pi}_t, \boldsymbol{\rho}_t^S \in \mathbb{R}_{++}^K$.

(2) $\Longrightarrow$ (1) [*Sufficiency*] By the definition of h-rationalisable, we only need to find support prices $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^K$ that satisfy the first-order conditions given by the HAD model.

Suppose there exist numbers $u_t \in \mathbb{R}$,$\lambda > 0$, $\{\delta^k\}_{k=1,...,K} \in (0,1]$, $\beta \in (0,1]$ and vectors $\mathbf{S}_t \in \mathbb{R}_+^K$, $\boldsymbol{\rho}_t^S \in \mathbb{R}_+^K$, $\boldsymbol{\rho}_t^q \in \mathbb{R}_{++}^K$, $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^K$ such that

$$u_s \leq u_t + \frac{1}{\beta^{t-1}}[\boldsymbol{\pi}_t'(\mathbf{q}_s - \mathbf{q}_t) + \boldsymbol{\rho}_t^{S'}(\mathbf{S}_s - \mathbf{S}_t)] \text{ for } \forall s, t = 2, ..., T \qquad [1]$$

$$\pi_t^k = \rho_t^{q,k} \text{ if } q_t^k > 0,\ t = 2, ..., T, k = 1, ..., K \qquad [2]$$

$$\rho_t^{S,k} = (1-\delta)[\rho_t^{q,k} - \lambda\rho_t^k] - [\rho_{t-1}^{q,k} - \lambda\rho_{t-1}^k] \text{ if } q_t^k > 0,\ t = 2, ..., T, k = 1, ..., K \qquad [3]$$

$$S_t^k = (1-\delta^k)S_{t-1}^k + q_{t-1}^k \text{ for } t = 2, ..., T, k = 1, ..., K \qquad [4]$$

Define $\tilde{\boldsymbol{\mu}}_t = \frac{1}{\beta^{t-1}}[\boldsymbol{\pi}_t, \boldsymbol{\rho}_t^S]'$ and $\tilde{\boldsymbol{\chi}}_t = [\mathbf{q}_t, \mathbf{S}_t]'$, so the inequality can be rewritten as

$$u_s \leq u_t + \tilde{\boldsymbol{\mu}}_t'(\mathbf{x}_s - \mathbf{x}_t) \text{ for } \forall s, t = 2, ..., T$$

which is a special case of the standard Afriat inequalities (where $\lambda_t = 1$ for $\forall t$). By Afriat's Theorem, the dataset $\mathbf{D} = \{\tilde{\boldsymbol{\mu}}_t, \tilde{\boldsymbol{\chi}}_t\}_{t=2,...,T}$ satisfies GARP and therefore there exists a locally nonsatiated, concave, continuous utility function $u(.)$ that is monotonic in $\tilde{\boldsymbol{\chi}}_t$ such that $\nabla u(\tilde{\boldsymbol{\chi}}_t) \leq \tilde{\boldsymbol{\mu}}_t$, with equality when $q_t^k > 0$ or $q_t^k = 0$ and $k \in \kappa_t$.

Either: 1. Suppose that if $q_t^k = 0$, then $k \in \kappa_t$ i.e. all observed zero purchases are due to rationing, so that $\nabla u(\tilde{\boldsymbol{\chi}}_t) = \tilde{\boldsymbol{\mu}}_t$ for $t = 2, ..., T$.

Recall that $\theta_t^k = \pi_t^k - \lambda\rho_t^k$ if $q_t^k > 0$. Given π_t^k, λ, $\rho_t^{S,k}$, and $\{\theta_t^k\}_{t=1,...,T}^{k \notin K_t}$, define $\theta_t^k = \sum_{i=1}^{j-1} \frac{1}{(1-\delta)^i}\rho_{t+1-i}^{S,k} + \frac{1}{(1-\delta)^j}(\rho_{t+1-j}^{S,k} + \theta_{t-j}^k)$ if $q_t^k = 0$ but if $q_{t-j}^k > 0$, where j is the number of periods since the good was last purchased. Substituting for θ_t^k into condition [3] and using the definition of $\rho_t^{S,k}$ gives the first-order condition for the stock variable:

$$\rho_t^{S,k} = \beta^{t-1}\frac{\partial u(\mathbf{q}_t, \mathbf{S}_t)}{\partial S_t^k} = (1-\delta^k)\theta_t^k - \theta_{t-1}^k \ ; t = 2, ..., T, k = 1, ..., K$$

Given π_t^k, λ, and θ_t^k, define $\rho_t^k = \frac{1}{\lambda}(\pi_t^k - \theta_t^k - \gamma_t^k)$ if $q_t^k = 0$, where $0 < \gamma_t^k < \pi_t^k - \theta_t^k$, so that $\rho_t^k = \begin{cases} \frac{1}{\lambda}(\pi_t^k - \theta_t^k) & q_t^k > 0 \\ \frac{1}{\lambda}(\pi_t^k - \theta_t^k - \gamma_t^k) & q_t^k = 0 \end{cases}$ and $\boldsymbol{\rho}_t \in \mathbb{R}_{++}^K$. Substituting for π_t^k obtains the first-order conditions for the HAD model:

$$\frac{\partial u(\mathbf{q}_t, \mathbf{S}_t)}{\partial q_t^k} = \frac{1}{\beta^{t-1}}[\lambda\rho_t^k + \theta_t^k] \text{ if } q_t^k > 0 \ ; t = 2, ..., T, k = 1, ..., K$$

$$\frac{\partial u(\mathbf{q}_t, \mathbf{S}_t)}{\partial q_t^k} = \frac{1}{\beta^{t-1}}[\lambda\rho_t^k + \theta_t^k + \gamma_t^k] \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t; t = 2, ..., T, k = 1, ..., K$$

Or: 2. Suppose that if $q_t^k = 0$, then $k \notin \kappa_t$ i.e. all observed zero purchases are due to corner solutions, so that $\nabla u(\tilde{\boldsymbol{\chi}}_t) \leq \tilde{\boldsymbol{\mu}}_t$, with equality when $q_t^k > 0$.

Recall that $\theta_t^k = \pi_t^k - \lambda\rho_t^k$ if $q_t^k > 0$. Given π_t^k, λ, $\rho_t^{S,k}$, and $\{\theta_t^k\}_{t=1,...,T}^{k \notin K_t}$, define $\theta_t^k = \sum_{i=1}^{j-1} \frac{1}{(1-\delta)^i}\rho_{t+1-i}^{S,k} + \frac{1}{(1-\delta)^j}(\rho_{t+1-k}^{S,k} + \theta_{t-j}^k)$ if $q_t^k = 0$ but if $q_{t-j}^k > 0$, where j is the number of periods since the good was last purchased. Substituting for θ_t^k into condition [3] and using the definition of $\rho_t^{S,k}$ gives the first-order condition for the stock variable:

$$\rho_t^{S,k} = \beta^{t}\frac{\partial u(\mathbf{q}_t, \mathbf{S}_t)}{\partial S_t^k} = (1-\delta^k)\theta_t^k - \theta_{t-1}^k \ ; t = 2, ..., T, k = 1, ..., K$$

Given π_t^k, λ, and θ_t^k, define $\rho_t^k = \frac{1}{\lambda}(\pi_t^k - \theta_t^k)$ if $q_t^k = 0$, so $\rho_t^k = \frac{1}{\lambda}(\pi_t^k - \theta_t^k)$ for $\forall q_t^k$. Substituting for π_t^k obtains the first-order conditions for the HAD model:'

$$\frac{\partial u(\mathbf{q}_t, \mathbf{S}_t)}{\partial q_t^k} = \frac{1}{\beta^{t-1}}[\lambda\rho_t^k + \theta_t^k] \text{ if } q_t^k > 0 \ ; t = 2, ..., T, k = 1, ..., K$$

$$\frac{\partial u(\mathbf{q}_t, \mathbf{S}_t)}{\partial q_t^k} \leq \frac{1}{\beta^{t-1}}[\lambda\rho_t^k + \theta_t^k] \text{ if } q_t^k > 0 \text{ and } k \notin \kappa_t \ ; t = 2, ..., T, k = 1, ..., K \ \square$$

Appendix 1B: Modelling storability as asset investment

Buying storable goods can be thought of as investing in assets. The asset investment model of Deaton (1992) can be interpreted in terms of storable goods. In this model, good k has a real "return" of $(1+r_t^k)$ in period t, where $r_t^k \in \mathbb{R}$ represents the real appreciation/depreciation in value, taken to be exogenous. Under the assumption of perfect foresight, the consumer knows all the real returns of all goods, but in general settings with uncertainty, the true r_t^k may not be known by the consumer at time $t-1$.

The budget constraint between time t and $t+1$ is $\sum_{k=1}^{K} \rho_k^t q_t^k = y_t + A_t$, where y_t is discounted income in period t (taken to be exogenous) and $A_t = \sum_{k=1}^{K}(1+r_t^k)\rho_{t-1}^k q_{t-1}^k$ is the current value of all "assets" owned at the start of period t i.e. the current value of storable goods purchased in $t-1$. The difference between this investment model and the investment model of Demuynck and Verriest (2013) is that assets appear in the budget constraint but are not included in the utility function. Thus the consumer does not get utility directly from the stock of storable goods, but can get utility indirectly by selling some of his stock to buy goods in the current period.

The consumer's lifetime maximisation problem is therefore:

$$\max_{q_t \in \mathbb{R}_+} \sum_{t=1}^{\infty} \beta^{t-1} u(\mathbf{q}_t) \text{ s.t. } \sum_{t=1}^{\infty} \sum_{k=1}^{K} \rho_t^k q_t^k = Y + \sum_{t=1}^{\infty} \sum_{k=1}^{K} (1+r_t^k)\rho_{t-1}^k q_{t-1}^k \text{ and}$$

$$q_t^k = 0 \text{ for any } k \in \kappa_t$$

where $Y = \sum_{t=1}^{\infty} y_t$ (total lifetime income, taken to be exogenous) and κ_t (the set of goods unavailable at time t) are generally not observed by the researcher.

The lagrangian is $\mathcal{L} = \sum_{t=1}^{\infty} \beta^{t-1} u(\mathbf{q}_t) + \lambda(Y + \sum_{t=1}^{\infty} \sum_{k=1}^{K} (1+r_t^k)\rho_{t-1}^k q_{t-1}^k) - \sum_{t=1}^{\infty} \sum_{k=1}^{K} \rho_t^k q_t^k) + \gamma_t^k (0 - q_t^k)$ and the first-order conditions are:

$$\beta^{t-1} \frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \lambda \rho_t^k \tilde{r}_{t+1}^k \text{ if } q_t > 0$$

$$\beta^{t-1} \frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} \leq \lambda \rho_t^k \tilde{r}_{t+1}^k \text{ if } q_t = 0 \text{ and } k \notin \kappa_t$$

$$\beta^{t-1}\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \lambda\rho_t^k\tilde{r}_{t+1}^k + \gamma_t^k \text{ if } q_t = 0 \text{ and } k \notin \kappa_t$$

for $t = 1, ..., T-1$, where $\tilde{r}_{t+1}^k = -r_{t+1}^k$. Using this notation, the first-order conditions are similar to those of the life cycle model in Section 1.2, except with the additional $\tilde{r}_{t+1}^k$ term.

Definition 1.5. A utility function $u(.)$ ($u : \mathbb{R}_+^K \to \mathbb{R}$) **d-rationalises** the dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ if there exist support prices $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^K$ (where $\pi_t^k = \rho_t^k$ if $q_t^k > 0$) and $\beta \in (0,1]$ such that for $\forall\ t \in \{1, ..., T\}$, $u(\mathbf{q}_t) \geq u(\mathbf{q})$ for any $\mathbf{q} \in \{\mathbf{q} \in \mathbb{R}_+^k : \boldsymbol{\pi}_t^{'}\mathbf{q} \leq \boldsymbol{\pi}_t^{'}\mathbf{q}_t\}$.

By Definition 1.5, the data $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ are therefore rationalisable by the Deaton (1992) asset investment model (d-rationalisable) if there exists a locally non-satiated, continuous, concave, and differentiable utility function $u(\mathbf{q}_t)$ and numbers $\lambda > 0$, $\gamma_t^k > 0$, $\tilde{r}_{t+1}^k \in \mathbb{R}$ and $\beta \in (0,1]$ that satisfy the first-order conditions given above.

The revealed preference conditions for the data $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ to be rationalisable by the Deaton (1992) asset investment model are given below.

Theorem 1.6 (Asset investment model). The following statements are equivalent:

(1) The data $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ are d-rationalisable.

(2) There exist numbers $u_t \in \mathbb{R}, \lambda > 0, \beta \in (0,1]$ and vectors $\tilde{\mathbf{r}}_{t+1} \in \mathbb{R}^K$, $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^K$ such that

$$u_s \leq u_t + \tfrac{1}{\beta^{t-1}}\lambda\tilde{\mathbf{r}}_{t+1}\boldsymbol{\pi}_t'(\mathbf{q}_s - \mathbf{q}_t) \text{ for } \forall s,t = 1, ..., T-1$$

$$\pi_t^k = \rho_t^k \text{ if } q_t^k > 0, \text{ for } \forall k = 1, ..., K; t = 1, ..., T-1\ .$$

***Proof*:**

[*Necessity*] The dataset $\mathbf{D} = \{i_t, \mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ is d-rationalisable if there exists a locally non-satiated, continuous, concave, and differentiable utility function $u(.)$ and numbers $\lambda > 0$, $\gamma_t^k > 0$, $\tilde{r}_{t+1}^k \in \mathbb{R}$, and $\beta \in (0,1]$ such that

$$\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \tfrac{1}{\beta^{t-1}}\lambda\rho_t^k\tilde{r}_{t+1}^k \text{ if } q_t^k > 0\ ;\ t = 1, ..., T-1, k = 1, ..., K$$

$$\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} \leq \tfrac{1}{\beta^{t-1}}\lambda\rho_t^k\tilde{r}_{t+1}^k \text{ if } q_t^k = 0 \text{ and } k \notin \kappa_t\ ;\ t = 1, ..., T-1, k = 1, ..., K$$

$$\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \tfrac{1}{\beta^{t-1}}[\lambda\rho_t^k\tilde{r}_{t+1}^k + \gamma_t^k] \text{ if } q_t^k = 0 \text{ and } k \in \kappa_t;\ t = 1, ..., T-1, k = 1, ..., K$$

By concavity of $u(.)$,

$$u_s \leq u_t + \nabla u(\mathbf{q}_t)'(\mathbf{q}_s - \mathbf{q}_t) \text{ for } \forall s, t = 1, ..., T-1$$

For $t = 1, ..., T-1, k = 1, ..., K$, define $\pi_t^k = \rho_t^k$ if $q_t^k > 0$ or $q_t^k = 0$ and $k \notin \kappa_t$, and $\pi_t^k = \rho_t^k + \frac{\gamma_t^k}{\lambda}$ if $q_t^k = 0$ and $k \in \kappa_t$. Note that $\pi_t^k > 0$ since ρ_t^k, γ_t^k, and $\lambda > 0$, so $\nabla u(\mathbf{q}_t) \leq \frac{1}{\beta^{t-1}} \lambda \tilde{\mathbf{r}}_{t+1} \boldsymbol{\pi}_t$, with equality if $q_t^k > 0$. Substituting $\frac{1}{\beta^{t-1}} \lambda \tilde{\mathbf{r}}_{t+1} \boldsymbol{\pi}_t$ for $\nabla u(\mathbf{q}_t)$ therefore preserves the inequality and gives statement (2):

$$u_s \leq u_t + \frac{1}{\beta^{t-1}} \lambda \tilde{\mathbf{r}}_{t+1} \boldsymbol{\pi}_t'(\mathbf{q}_s - \mathbf{q}_t) \text{ for } \forall s, t = 1, ..., T-1$$

$$\pi_t^k = \rho_t^k \text{ for any } q_t^k > 0 \text{ for } \forall k = 1, ...K, t = 1, ..., T-1$$

where $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^K$.

(2) $\implies$ (1) [*Sufficiency*] By the definition of d-rationalisable, we only need to find support prices $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^K$ that satisfy the first-order conditions given by the asset invesment model. Suppose there exist numbers $u_t \in \mathbb{R}$, $\lambda > 0$, $\beta \in (0, 1]$, and vectors $\tilde{\mathbf{r}}_{t+1} \in \mathbb{R}^K$, $\boldsymbol{\pi}_t \in \mathbb{R}_{++}^K$ such that

$$u_s \leq u_t + \frac{1}{\beta^{t-1}} [\lambda \tilde{\mathbf{r}}_{t+1} \boldsymbol{\pi}_t'(\mathbf{q}_s - \mathbf{q}_t)] \text{ for } \forall s, t = 1, ..., T-1$$

$$\pi_t^k = \rho_t^k \text{ for any } q_t^k > 0 \text{ for } \forall k = 1, ...K, t = 1, ..., T-1$$

Define $\hat{\boldsymbol{\mu}}_t = \frac{1}{\beta^{t-1}} \tilde{\mathbf{r}}_{t+1} \boldsymbol{\pi}_t$, so the inequality can be rewritten as

$$u_s \leq u_t + \lambda \hat{\boldsymbol{\mu}}_t'(\mathbf{q}_s - \mathbf{q}_t) \text{ for } \forall s, t = 1, ..., T-1$$

which is a special case of the standard Afriat inequalities (where $\lambda_t = \lambda$ for $\forall t$). By Afriat's Theorem, the dataset $\mathbf{D'} = \{\hat{\boldsymbol{\mu}}_t, \mathbf{q}_t\}_{t=1,...,T}$ satisfies GARP and therefore there exists a locally nonsatiated, concave, continuous utility function $u(.)$ that is monotonic in $\mathbf{q}_t$ such that $\nabla u(\mathbf{q}_t) \leq \lambda \hat{\boldsymbol{\mu}}_t$, with equality when $q_t^k > 0$ or $q_t^k = 0$ and $k \in \kappa_t$.

Either: 1. Suppose that if $q_t^k = 0$, then $k \in \kappa_t$ i.e. all observed zero purchases are due to rationing, so that $\nabla u(\mathbf{q}_t) = \frac{1}{\beta^{t-1}} \lambda \tilde{\mathbf{r}}_{t+1} \boldsymbol{\pi}_t$ for $t = 1, ..., T-1$.

Given π_t^k and λ, define $\rho_t^k = \pi_t^k - \frac{\gamma_t^k}{\lambda}$ if $q_t^k = 0$, where $0 < \gamma_t^k < \lambda\pi_t^k$, so that $\rho_t^k = \begin{cases} \pi_t^k & q_t^k > 0 \\ \pi_t^k - \frac{\gamma_t^k}{\lambda} & q_t^k = 0 \end{cases}$ and $\boldsymbol{\rho}_t \in \mathbb{R}_{++}^K$. Substituting for π_t^k obtains the first-order conditions for the asset investment model:

$$\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \frac{1}{\beta^{t-1}}\lambda\rho_t^k\tilde{r}_{t+1}^k \text{ if } q_t^k > 0 \ ; t = 1, ..., T-1, k = 1, ..., K$$
$$\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \frac{1}{\beta^{t-1}}[\lambda\rho_t^k\tilde{r}_{t+1}^k + \gamma_t^k] \text{ if } q_t^k = 0; t = 1, ..., T-1, k = 1, ..., K$$

Or: 2. Suppose that if $q_t^k = 0$, then $k \notin \kappa_t$ i.e. all observed zero purchases are due to corner solutions, so that $\nabla u(\mathbf{q}_t) \leq \frac{1}{\beta^{t-1}}\lambda\tilde{\mathbf{r}}_{t+1}\boldsymbol{\pi}_t$, with equality when $q_t^k > 0$.

Given π_t^k, define $\rho_t^k = \pi_t^k$ for $q_t^k = 0$, so $\rho_t^k = \pi_t^k$ for $\forall q_t^k$. Substituting for π_t^k obtains the first-order conditions for the asset investment model:

$$\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} = \frac{1}{\beta^{t-1}}\lambda\rho_t^k\tilde{r}_{t+1}^k \text{ if } q_t^k > 0 \ ; t = 1, ..., T-1, k = 1, ..., K$$
$$\frac{\partial u(\mathbf{q}_t)}{\partial q_t^k} \leq \frac{1}{\beta^{t-1}}\lambda\rho_t^k\tilde{r}_{t+1}^k \text{ if } q_t^k = 0 \ ; t = 1, ..., T-1, k = 1, ..., K \ \square$$

The proof is analogous to that of the life cycle model, with an additional term for the real appreciation/depreciation in value of each good, $\tilde{r}_{t+1}^k$. These inequalities are nonlinear in unknowns $(\beta, \lambda, \tilde{r}_{t+1}^k)$. If the researcher is able to observe $\{\tilde{r}_{t+1}^k\}$, then the solution to these inequalities can be solved by via a grid search on β as in Section 1.2. Otherwise, a solution has to be found using non-linear programming methods.

As with the life cycle model, the discounted prices ρ_t can be constructed from the numbers $\{\pi_t^k\}_{t=1,...,T}^{k=1,...,K}$ and λ that satisfy the inequalities (obtained from the solution to the revealed preference test). For nonzero purchases and for zero purchases due to corner solutions, the discounted price equals the support price ($\pi_t^k = \rho_t^k$). For zero purchases due to rationing, the discounted price is a function of the support price π_t^k, λ (the marginal utility of discounted total lifetime expenditure), and γ_t^k. γ_t^k represents the marginal cost to utility from rationing, which is unobservable to the researcher but can be any number such that $0 < \gamma_t^k < \lambda\pi_t^k$, so it is possible for a set of values of γ_t^k to rationalise the data for a given solution to the inequalities.

2. Happiness and Economics

2.1 Motivation

Happiness measures and other subjective evaluations of well-being have become increasingly important in policy debates. Over the past decade there have been national and international initiatives aimed at using these measures of well-being and societal progress to inform policy-making, rather than solely relying on economic indicators such as GDP[36]. The academic literature has also applied happiness measures to economic settings, for both positive and normative purposes. Some studies estimate 'happiness functions', which are exploratory regressions aimed at finding statistically significant determinants of individual or country-level happiness[37]. Other studies actually use self-reported happiness as a proxy for utility or individual welfare, taking the estimates from these 'happiness functions' to represent an individual's preferences over the covariates. This approach has been used to value non-market goods such as health and air quality, but also to conduct welfare analysis, including calculat-

[36]For example, the Office for National Statistics recently started using subjective measures to assess economic well-being, such as using household surveys of satisfaction with income to evaluate their financial situation (see http://www.ons.gov.uk/peoplepopulationandcommunity/wellbeing/bulletins/measuringnationalwellbeing/2015-09-23). The OECD also launched a project on better measures of living conditions and societal well-being, along with guides for measuring well-being (see https://www.oecd.org/statistics/measuring-well-being-and-progress.htm). An overview and further discussion of alternative indicators of well-being and progress can be found in Schepelmann et al (2010).

[37]For example, Oswald (1997) and Frank (1997) explore the relationship between country-level happiness indices and macroeconomic variables such as per capita income or GDP. Ferrer-i-Carbonell and Frijters (2004) discuss how to use econometric techniques such as linear or ordered probit/logit regressions to estimate a 'happiness function' based on happiness ratings and individual characteristics such as employment status, income, and demographics.

ing compensating and equivalent variation[38].

Although happiness measures can be useful when considered alongside economic variables, there are some issues to consider when using them as proxies of utility or individual welfare. One concern with using happiness ratings in economic settings is the subjectivity of these statements. Kahnemann and Krueger (2006) outline various factors that do not influence the measurement of economic variables but may affect the reliability of happiness reports, such as the context and framing of survey questions, or individual-specific memories and experiences. A key assumption when using happiness as a proxy for utility is that the former, like the latter, represents a rational preference over some domain, so that observed behaviour in that domain can be considered as 'happiness-maximising'. Although studies in psychology have found that self-rated happiness is consistent with observable behaviour that is generally accepted as indicative of well-being[39], the consistency between preferences implied by happiness measures and utility maximisation has not been verified in many non-experimental economic settings.

Only a few studies to date have examined the relationship between these two sets of preferences in an economic context without using a controlled experiment. The study most closely related to this chapter is Bernal et al (LISS WP, 2014), which looks at the correlation between happiness measures and utility using Dutch household survey data. Utility levels were estimated from a model of collective household behaviour, in which all household resources

[38]Specifically, the happiness or satisfaction of individual i at time t (H_{it}) is assumed to be a function of income (x_{it}) and other covariates ($\mathbf{z}$):

$$H_{it} = \varphi log(x_{it}) + \sum_k \Lambda_k z_k + \epsilon_{it}$$

The ratio of coefficients, $-\frac{\Lambda_k}{\varphi}$, is then interpreted as an individual's marginal rate of substitution between income and covariate k. Table 2.23 presents examples of studies from a variety of fields in economics that use this approach or a similar approach to conduct welfare analysis or estimate the value of a non-market good. For a more comprehensive overview of happiness measures and their economic applications, see Ferrer-i-Carbonell (2013), Frey and Stutzer (2012), and MacKerron (2012).

[39]Happiness reports are shown to have high consistency, reliability, and stability over time (Headey and Wearing, 1989; Frey and Stutzer, 2002). Self-rated happiness was also found to be predictive of observable behaviour such as active social relationships and regular exercise, which is generally accepted as indicative of well-being (Frank, 1997).

are distributed among members and each household member receives utility according to their private consumption of goods and leisure time. These estimated utility levels were not significantly correlated with reported happiness ratings of general happiness or life satisfaction, though the authors note that this imperfect correlation may be due to the definition of leisure used[40], assumptions about the utility function[41], or factors at the time of the survey that affected individual responses to the happiness questions[42].

Bernal et al's approach can be described as structural because it relies on a specific model of human behaviour to explain the data. Their assumption that households behaved according to the model specified was required to estimate unobserved utility levels, but was not verified. Therefore, another explanation for their results is that their model of household behaviour was incorrect. Also, their use of linear regression to measure the correlation between happiness and utility requires strong assumptions about the relationship between these two variables. Since utility levels determine an individual's preference ranking over options rather than the magnitude of an individual's preference for one option over another, a definitive check for consistency would only use preference rankings. The non-parametric revealed preference approach is one way to resolve these issues. Consistency between happiness measures and utility maximisation can be tested for without assuming cardinality, and assumptions about the particular decision-making process can also be verified.

This chapter will use a nonparametric revealed preference approach to check for consistency between happiness measures and preferences over consumption, using household survey data. Since expenditure data is at the household level but happiness measures are reported

[40]Bernal et al (2014) found that 'pure' leisure (time spent on leisure activities) was more closely related with happiness measures compared to 'constructed leisure' (hours worked minus time spent sleeping and personal care, a definition usually applied in labour supply models when detailed time use information is unavailable).

[41]Bernal et al (2014) assumed that utility was a function of private consumption (total private expenditure on all goods) and leisure, but there may be other variables that were not accounted for, such as public goods consumption. Also, their estimates of utility are based on the Linear Expenditure System functional form.

[42]For example, the context or framing of survey questions, or unexpected events that influenced the individual's mood at the time of survey.

at the individual level, revealed preference tests will be conducted under different assumptions on household behaviour[43] and different measures of happiness[44]. Consistency means that a household's behaviour satisfies both the assumptions of the model and the preference ordering given by their happiness ratings. Finding such consistency would support the current use of happiness measures in normative economics.

The rest of the chapter is organised as follows: Section 2.2 develops the key concepts in revealed preference theory mentioned in the Introduction and outlines the specific models used. Section 2.3 describes the data used, with a particular focus on the happiness measures. Section 2.4 presents the results of the revealed preference tests and discuss some reasons why consistency may not hold. Section 2.5 concludes.

2.2 Revealed preference theory

2.2.1 Unitary household behaviour

Afriat's Theorem (stated in the Introduction) gave necessary and sufficient conditions for an individual's behaviour to be consistent with utility maximisation. Revealed preference tests for individual data be applied to household-level data by assuming that all members in a household have the same preferences (referred to as 'unitary household behaviour'), so the household acts as if it was a single decision maker.

Since Afriat's Theorem deals with individual behaviour that is independent from the behaviour of others, the GARP test will be run separately for each household, using only data related to that household. There are therefore no restrictions on the heterogeneity in pref-

[43] Under the unitary approach (household members all have the same preferences), the revealed preference tests for individuals described in the Introduction apply. Under the collective approach (household members can have different preferences), models of collective household behaviour as in Cherchye (2011b) are used instead, which take into account each individual's contribution to the final decision.

[44] Overall and domain-specific happiness measures are reported by the respondent and his/her spouse.

erences across households. To simplify notation, the household subscript will be omitted from the subsequent discussion, and data for each household can be thought of as a separate dataset.

As mentioned in the Introduction, if a household satisfies GARP, the numbers $\{y_{st}\}_{s,t=1,...,T} \in \{0,1\}$ that solve the set of linear inequalities given by (3) in Theorem 0.1 can be interpreted as a preference ordering over the bundles purchased[45]. If happiness measures are consistent with revealed preferences, then the preference ordering given by happiness ratings should be consistent with the preferences implied by the observed choices. That is, if a household is reportedly happier at time A than time B, then consumption at A must cost at least as much as consumption at B (as in condition (3) of Definition 0.1).

Therefore, if the preferences implied by happiness measures and utility maximisation are consistent, then the happiness ratings should yield a solution to the inequalities given in Afriat's Theorem. Consistency for a given household can therefore be verified by constructing $\{y_{st}\}_{s,t=1,...,T}$ according to the ordering given by their happiness numbers $H_1, ..., H_T$, and checking whether the inequalities are satisfied (the GARP test for the unitary model that uses happiness numbers is referred to as GARP-H)[46]. Figure 2.1 shows examples of consistency and inconsistency between preferences implied by consumption behaviour and happiness ratings. The inconsistency arises because the household is happier with the cheaper bundle B, which violates condition (3) of GARP.

[45]Since the solution to the set of linear inequalities may not be unique, there could be multiple preference orderings that are consistent with utility-maximisation.

[46]For example, if $H_1 = 5$ and $H_2 = 7$, then $y_{12} = 0$ and $y_{21} = 1$.

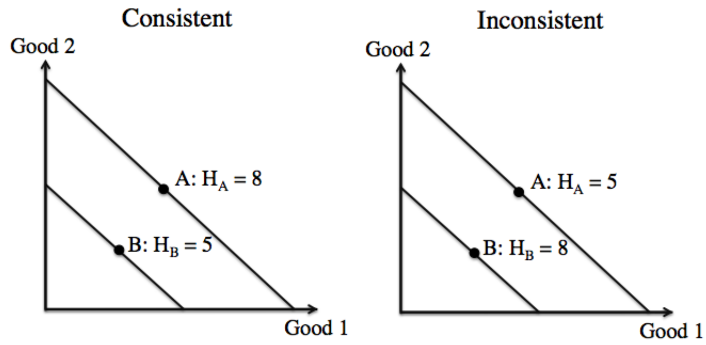

Figure 2.1: Examples of consistency and inconsistency

Figure 2.2 demonstrates why happiness measures can be useful in economic settings. As mentioned in the Introduction, it is not always possible to construct a complete and unique preference ordering over consumption bundles using only price and quantity information. For example, bundles A and B in Figure 2.2 satisfy GARP, but there is no way to determine whether bundle A is preferred to bundle B (it is possible to find preferences such that A is preferred to B but also vice versa). However, if the data is consistent with GARP-H, then the happiness numbers can help form a complete preference ordering between periods by enabling comparisons between all bundles. In this case, the happiness ratings indicate that A is preferred to B.

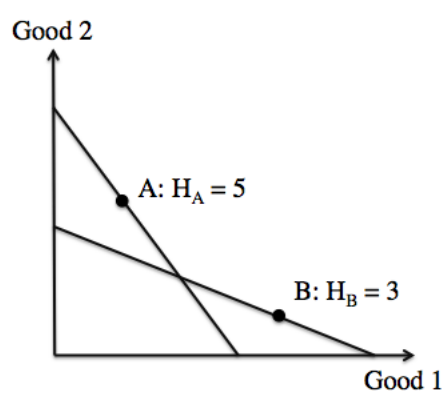

Figure 2.2: Constructing preferences from happiness numbers

2.2.2 Collective household behaviour

In the unitary approach, households were treated as if they acted as a single decision maker, in order to apply the revealed preference results of Afriat (1967) and Varian (1982). While this approach makes no difference for single-person households, previous studies have found it is not a good fit for multi-person households[47], which challenges the assumptions of how individual preferences are aggregated at the household level.

In the collective approach, each household member can have their own preferences[48]. All the goods that a household purchases in a given period is divided between household members through a 'sharing rule', which represents the bargaining power of each member. There are many variations of collective household behaviour[49], but this chapter uses the standard model of Cherchye et al (2011b). As in the unitary case, the collective approach of Cherchye et al (2011b) is nonparametric and so does not impose any parametric structure on individual preferences.

In the collective household model of Cherchye et al (2011b), there are K private goods and N public goods, which are divided between M household members. Household consumption in each period $t = 1, ..., T$ consists of both private goods ($\mathbf{q}_t \in \mathbb{R}^K_+$) and public goods ($\mathbf{Q}_t \in \mathbb{R}^N_+$). Private goods can be used by one individual only and have no consumption externalities (for example, clothing and medicine), whereas public goods either have consumption externalities or are non-excludable (for example, an internet connection or a TV).

For private goods, the total amount purchased in each period is divided among the members and each member pays the observed price $\mathbf{p}_t \in \mathbb{R}^K_{++}$ per unit consumed. The bundle of private goods that member m receives in period t is denoted as $\bar{\mathbf{q}}^m_t$.

[47] For example, see Browning and Chiappori (1998), Cherchye and Vermeulen (2008), Cherchye et al (2009), and Cherchye et al (2011c).

[48] With this setup, the unitary model is a special case of the collective model in which each household member has the same preferences.

[49] See, for example, Cherchye et al (2011c), and Cherchye et al (2015).

For public goods, each member consumes the aggregate amount $\mathbf{Q}_t$ but pays a non-zero fraction of the observed price according to how much he or she values the good (a similar concept to Lindahl pricing). The amount paid by member m for public goods in period t (also referred to as **feasible personalised prices**) is denoted as $\bar{\mathbf{P}}_t^m$.

The total amounts of private and public goods purchased ($\mathbf{q}_t$ and $\mathbf{Q}_t$) is observable, as well as the market prices paid for these goods ($\mathbf{p}_t$ and $\mathbf{P}_t$ respectively)[50]. The observable variables place some restrictions on the unobservable variables: feasible personalised prices must add up to the observed public goods prices ($\sum_{m=1}^{M} \bar{\mathbf{P}}_t^m = \mathbf{P}_t$), and member-specific private consumption must add up to the observed total amounts purchased ($\sum_{m=1}^{M} \bar{\mathbf{q}}_t^m = \mathbf{q}_t$). The **income share** received by each member (denoted η_t^m for member m) is usually unobserved, but must be equal to that member's total expenditure on public and private consumption: $\eta_t^m = \mathbf{p}_t \cdot \bar{\mathbf{q}}_t^m + \bar{\mathbf{P}}_t^m \cdot \mathbf{Q}_t$.

Some household surveys ask respondents to report either the amount or share of a subset $k \subseteq K$ of goods that they privately consumed, so $\bar{\mathbf{q}}_t^m$ can be partially observed. These private goods are referred to as '**assignable**' and are denoted $q_t^{A,m,n}$ for member m and good n. This additional information places further restrictions on private consumption: the actual amount consumed must be greater than or equal than the reported amount ($\bar{q}_t^{m,n} \geq q_t^{A,m,n}$ for $k \subseteq K$, and $\bar{q}_t^{m,n} \geq 0$ otherwise). **Feasible personalized quantities** ($\bar{\mathbf{q}}_t \in \mathbb{R}_+^K$) are therefore defined as any division of aggregate private good consumption that satisfies the adding-up constraint and restrictions on assignable goods (that is, $\bar{\mathbf{q}}_t^m \geq \mathbf{q}_t^{A,m}$ for $m = 1, ..., M$ and $\sum_{m=1}^{M} \bar{\mathbf{q}}_t^m = \mathbf{q}_t$). Figure 2.3 summarises the terms and relationships used in this collective model.

[50] Regarding notation: throughout this chapter, a bar is used to distinguish between observed and unobserved variables. Variables related to private goods are in lower case, variables related to public goods are in upper case. In both cases, bold font refers to vectors and non-bold font refers to scalars. As in Section 2.2.1, to simplify notation, the household subscript is omitted, and data for each household can be thought of as a separate dataset.

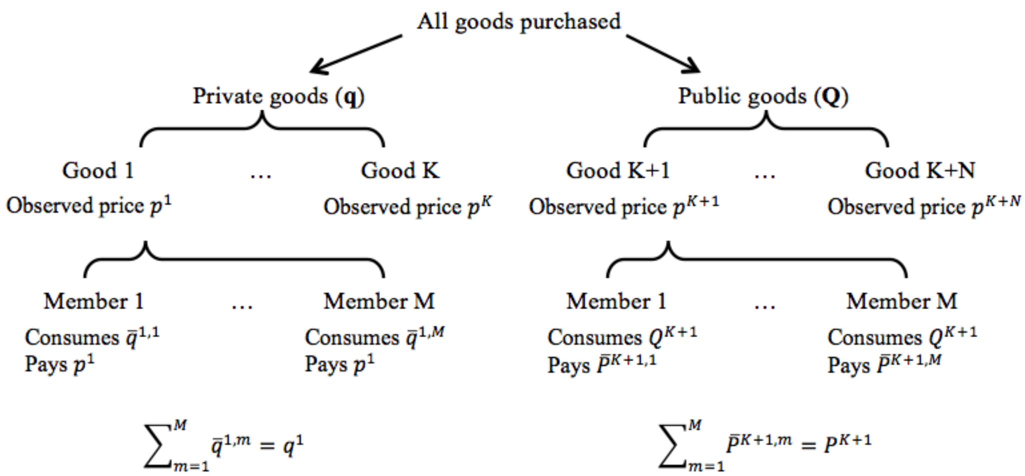

Figure 2.3: Collective model of household behaviour

As shown in Figure 2.3, observed household behaviour can be thought of as the result of a two-step allocation procedure: the sharing rule is used to divide total household income among members, then each individual chooses private and public good consumption to maximise his or her own preferences. Individual preferences are given by the utility functions $U^m(\bar{\mathbf{q}}^m, \mathbf{Q})$ for $m = 1, ..., M$, which are assumed to be well-behaved (non-satiated, continuous, concave, and monotonic in $\bar{\mathbf{q}}^m$ and $\mathbf{Q}$). As formally defined below, for household data to be consistent with the collective model, there must be a set of weakly positive feasible personalised (private) quantities, strictly positive (public) prices, and income shares such that the observed household behaviour is consistent with individual-level maximisation:

Definition 2.1.[51] A combination of utility functions $U^1, ..., U^M$ **collectively rationalises** (could have generated) the data $\mathbf{D} = \{\mathbf{p}_t, \mathbf{q}_t, \mathbf{P}_t, \mathbf{Q}_t\}_{t=1,...,T}$ if for all t, m there exist feasible personalised prices ($\bar{\mathbf{P}}_t \in \mathbb{R}^N_{++}$)and quantities ($\bar{\mathbf{q}}_t \in \mathbb{R}^K_+$) such that $U^m(\mathbf{q}^m_t, \mathbf{Q}_t) \geq U^m(\hat{\mathbf{q}}^m_t, \hat{\mathbf{Q}})$ for all $\hat{\mathbf{q}}^m_t \in \mathbb{R}^K_+$ and $\hat{\mathbf{Q}} \in \mathbb{R}^N_+$ such that $\mathbf{p}_t \cdot \bar{\mathbf{q}}^m_t + \bar{\mathbf{P}}^m_t \cdot \mathbf{Q}_t \geq \mathbf{p}_t \cdot \hat{\mathbf{q}}^m_t + \bar{\mathbf{P}}^m_t \cdot \hat{\mathbf{Q}}$.

This definition implies that consistency with the collective model requires individual behaviour to be consistent with utility maximisation. As mentioned in the Introduction, GARP

[51]Paraphrased from Cherchye et al (2011b).

is a necessary and sufficient condition for individual utility maximisation given the observed prices and quantities (Varian, 1982). The definition of GARP for the collective model is a generalisation of GARP for the unitary model, and each condition has the same interpretation:

Definition 2.2. The set of feasible personalised prices and quantities $\{\mathbf{p}_t, \bar{\mathbf{q}}_t^m, \bar{\mathbf{P}}_t^m, \mathbf{Q}_t\}_{t=1,...,T}^{m=1,...,M}$ for the dataset **D satisfies GARP** if and only if we can construct preference relations R_m and R_m^0 such that for all $r, s, t, u, v \in T$:

1. if $\mathbf{p}_t \cdot \bar{\mathbf{q}}_t^m + \bar{\mathbf{P}}_t^m \cdot \mathbf{Q}_t \geq \mathbf{p}_t \cdot \mathbf{q}_s^m + \bar{\mathbf{P}}_t^m \cdot \mathbf{Q}_s$, then $(\bar{\mathbf{q}}_t^m, \mathbf{Q}_t) R_m^0 (\bar{\mathbf{q}}_s^m, \mathbf{Q}_s)$

2. if $(\bar{\mathbf{q}}_r^m, \mathbf{Q}_r) R_m^0 (\bar{\mathbf{q}}_s^m, \mathbf{Q}_s)$, $(\bar{\mathbf{q}}_s^m, \mathbf{Q}_s) R_m^0 (\bar{\mathbf{q}}_t^m, \mathbf{Q}_t)$, ..., $(\bar{\mathbf{q}}_u^m, \mathbf{Q}_u) R_m^0 (\bar{\mathbf{q}}_v^m, \mathbf{Q}_v)$, then $(\bar{\mathbf{q}}_r^m, \mathbf{Q}_r) R_m (\bar{\mathbf{q}}_v^m, \mathbf{Q}_v)$

3. if $(\bar{\mathbf{q}}_t^m, \mathbf{Q}_t) R_m (\bar{\mathbf{q}}_s^m, \mathbf{Q}_s)$, then $\mathbf{p}_s \cdot \bar{\mathbf{q}}_t^m + \bar{\mathbf{P}}_s^m \cdot \mathbf{Q}_t \geq \mathbf{p}_s \cdot \mathbf{q}_s^m + \bar{\mathbf{P}}_s^m \cdot \mathbf{Q}_s$.

Since some prices and quantities are not observable, it is not possible to check for GARP directly. However, Cherchye et al (2011b) derive necessary and sufficient conditions for data to be consistent with the collective model when only aggregate household expenditure on each good and total household expenditure in each period (defined as $x_t = \mathbf{p}_t \cdot \mathbf{q}_t + \mathbf{P}_t \cdot \mathbf{Q}_t$) is observed. The revealed preference test for the collective model follows from this definition of GARP.

Theorem 2.1.[52] Given the dataset $\mathbf{D} = \{\mathbf{p}_t, \mathbf{q}_t, \mathbf{P}_t, \mathbf{Q}_t, x_t\}_{t=1,...,T}$, the following conditions are equivalent:

(1) The dataset **D** is consistent with utility maximisation according to the collective model of household behaviour, and therefore there exists M well-behaved utility functions $U^1, ..., U^M$ that could have generated the data.

(2) There exist feasible personalised prices and quantities $\{\mathbf{p}_t, \bar{\mathbf{q}}_t^m, \bar{\mathbf{P}}_t^m, \mathbf{Q}_t\}_{t=1,...,T}^{m=1,...,M}$ that sat-

[52] Paraphrased from Proposition 1 and 2, Cherchye et al (2011b). See Cherchye et al (2011b) for proof.

isfy GARP.

(3) There exist $\bar{\mathbf{P}}_t^m \in \mathbb{R}_+^N$, $\bar{\mathbf{q}}_t^m \in \mathbb{R}_+^K$, $\eta_t^m \in \mathbb{R}_+$, and $y_{st}^m \in \{0,1\}$ such that

(I) $\sum_{m=1}^M \bar{\mathbf{P}}_t^m = \mathbf{P}_t$

(II) $\bar{\mathbf{q}}_t^m \geq \mathbf{q}_t^{A,m}$ and $\sum_{m=1}^M \bar{\mathbf{q}}_t^m = \mathbf{q}_t$

(III) $\eta_t^m = \mathbf{p}_t \cdot \bar{\mathbf{q}}_t^m + \bar{\mathbf{P}}_t^m \cdot \mathbf{Q}_t$

(IV) $\eta_s^m - (\mathbf{p}_s \cdot \bar{\mathbf{q}}_t^m + \bar{\mathbf{P}}_s^m \cdot \mathbf{Q}_t) < x_s y_{st}^m$

(V) $y_{su}^m + y_{ut}^m \leq 1 + y_{st}^m$

(VI) $\eta_t^m - (\mathbf{p}_t \cdot \bar{\mathbf{q}}_s^m + \bar{\mathbf{P}}_t^m \cdot \mathbf{Q}_s) < x_t(1 - y_{st}^m)$

for $s, t \in 1, ..., T$ and $m = 1, ..., M$, where $y_{st}^m = 1$ if $(\bar{\mathbf{q}}_s^m, \mathbf{Q}_s) R_m (\bar{\mathbf{q}}_t^m, \mathbf{Q}_t)$, given the feasible personalised prices and quantities.

Condition (3) is a mixed-integer programming problem that is the collective equivalent to the integer programming problem of the unitary model, with additional conditions on the personalised prices, quantities, and income shares[53]. Constraints (I), (II), and (III) come directly from the definition of feasible personalised prices, quantities, and income shares respectively, while constraints (IV) to (VI) are based on the GARP conditions (1)-(3) given in Definition 2.2.

To see this relationship more clearly, the left-hand side of condition (IV) can be rewritten as $(\mathbf{p}_s \cdot \bar{\mathbf{q}}_s^m + \bar{\mathbf{P}}_s^m \cdot \mathbf{Q}_s) - (\mathbf{p}_s \cdot \bar{\mathbf{q}}_t^m + \bar{\mathbf{P}}_s^m \cdot \mathbf{Q}_t)$. If this expression is non-negative (i.e. $(\mathbf{p}_s \cdot \bar{\mathbf{q}}_s^m + \bar{\mathbf{P}}_s^m \cdot \mathbf{Q}_s) \geq (\mathbf{p}_s \cdot \bar{\mathbf{q}}_t^m + \bar{\mathbf{P}}_s^m \cdot \mathbf{Q}_t)$), then by condition (1) of Definition 2.2, $(\bar{\mathbf{q}}_s^m, \mathbf{Q}_s) R_m^0 (\bar{\mathbf{q}}_t^m, \mathbf{Q}_t)$ and therefore $y_{st}^m = 1$. In this case, constraint (IV) restricts the difference between individual m's expenditure on the revealed preferred bundle and expenditure on the revealed worse bundle to be strictly less than total household expenditure[54]. If

[53] In a mixed-integer programming problem, some variables are constrained to only take on integer values. In this case, it is the binary variables $\{y_{st}^m\}_{t=1,...,T}^{m=1,...,M}$, which represent the preference relationships between time periods.

[54] Since feasible personalised quantities and income shares can be zero, this constraint accounts for the special case in which the entire household income is given to one individual.

$(\mathbf{p}_s \cdot \bar{\mathbf{q}}_s^m + \bar{\mathbf{P}}_s^m \cdot \mathbf{Q}_s) < (\mathbf{p}_s \cdot \bar{\mathbf{q}}_t^m + \bar{\mathbf{P}}_s^m \cdot \mathbf{Q}_t))$, then $y_{st}^m = 0$ and individual m's expenditure on the revealed preferred bundle (period t instead of period s this case) is required to be strictly less than zero[55]. By similar reasoning, constraint (VI) is based on condition (3) of Definition 2.2 (note the switching of s and t subscripts for the left-hand side and for total household income compared to constraint (IV), hence the relationship with y_{st}^m is also reversed). Constraint (V) imposes transitivity[56].

For household behaviour to be consistent with the collective model, there must be a solution to the mixed integer linear programming problem given by condition (3). If a solution exists, then by condition (3), these numbers correspond to the feasible personalised prices, quantities, and income shares of each household member, and the binary variables $\{y_{st}^m\}_{s,t=1,\ldots,T}^{m=1,\ldots,M}$ represent the between-period preference relationships of each household member (as in the unitary model)[57]. The GARP test for the collective model that uses happiness numbers will be referred to as GARP-CH.

The binary variables $\{y_{st}^m\}_{t=1,\ldots,T}^{m=1,\ldots,M}$ depend on the unobserved values of feasible prices and quantities, which makes the programming problem quite complicated, especially if the number of household members and/or time periods is large. To simplify the programming problem, the revealed preference test was run conditioning on the values of $\{y_{st}^m\}_{t=1,\ldots,T}^{m=1,\ldots,M}$. First, all possible rational combinations of $\{y_{st}^m\}_{t=1,\ldots,T}^{m=1,\ldots,M}$ were specified[58], and each of these combinations was substituted in place of $\{y_{st}^m\}_{t=1,\ldots,T}^{m=1,\ldots,M}$ until a solution was found (data passed) or all combinations were exhausted (data failed).

[55]The intuition being that if period t's bundle is revealed preferred to period s' bundle, then by optimality of the chosen bundle in period s given individual m's income share, the period t bundle should have been unaffordable in period s.

[56]By condition 2 (Definition 2.2), if $(\bar{\mathbf{q}}_s^m, \mathbf{Q}_s) R_m^0 (\bar{\mathbf{q}}_u^m, \mathbf{Q}_u)$ i.e. $y_{su}^m = 1$, $(\bar{\mathbf{q}}_u^m, \mathbf{Q}_u) R_m^0 (\bar{\mathbf{q}}_t^m, \mathbf{Q}_t)$ i.e. $y_{ut}^m = 1$, then $(\bar{\mathbf{q}}_s^m, \mathbf{Q}_s) R_m^0 (\bar{\mathbf{q}}_t^m, \mathbf{Q}_t)$ i.e. y_{st}^m must equal 1.

[57]Note that again it is possible for many solutions to exist.

[58]These are combinations of 0 and 1 that satisfy constraint (V) in Theorem 2.1, and only depend on the number of time periods and members of each household.

2.2.3 The usefulness of stated preferences

In this context, using happiness numbers can both significantly reduce the computational burden and impose more structure on the data. Since happiness numbers are real, non-negative integers, they represent a complete and transitive preference relation. As mentioned in Section 2.2.1, happiness responses can be used to fully specify each individual's preferences across time periods. The revealed preference test would then involve checking for the existence of feasible personalised prices, quantities, and income shares that are consistent with the preference relations given by the happiness numbers. As with the unitary approach in Section 2.2.1, only the relationship between individual-specific numbers is important. Therefore, only the assumption of within-person consistency is required[59], so that an increase in rated happiness from 6 in 2011 to 7 in 2012 means that the respondent prefers his consumed bundle in 2012 to his consumed bundle in 2011. This assumption is crucial for happiness numbers to be meaningful in welfare analysis, but may not hold if behaviour is non-neoclassical[60].

Figure 2.4 illustrates how an individual's happiness numbers $\{H_k\}_{k=1,\dots,T}$ can be used to define their complete preference ordering and hence y_{st} between periods for $k = 1, 2, 3$. For example, $y_{21} = 1$ since $8 > 5$. Note that even using the unitary approach, it is not always possible to construct a complete preference ordering between all periods (see Figure 2.2). Imposing the complete preference ordering using happiness numbers may reduce the pass rate if the additional restrictions implied are not consistent with observed behaviour. The following section will investigate whether or not happiness numbers represent a preference ordering that is consistent with individual utility maximisation according to the collective

[59] An individual's perception of happiness ratings should not change over time, so that rating of 6 in 2011 refers to the same standard of well-being as a rating of 6 in 2012 or 2013.

[60] For example, in models of reference dependence, individual utility is measured with respect to a reference point. It is possible for utility levels to not correspond to absolute expenditure if the reference point changes e.g. if an individual has an unexpected pay-rise in 2011 but receives the same (higher) pay in 2012 and 2013, it is possible for him to have higher utility in 2011 compared to 2012 and 2013 even though his spending power remains the same.

model.

Period (k)	H_k
1	5
2	8
3	6

		t		
	y_{st}	1	2	3
s	1	1	0	0
	2	1	1	1
	3	1	0	1

Figure 2.4: Constructing preferences from happiness numbers

2.3 Data

The Japan Household Panel Survey (JHPS) started in 2009 and is conducted on the 31st of January every year. The purpose of the JHPS is to obtain comprehensive information about common research topics such as household composition, expenditure, employment, health, and housing situation from a representative subset of the Japanese population. The JHPS is still ongoing, but the 2013 survey was the latest data available to researchers at the time of this study. The dataset consists of 16,054 observations from 4022 households. The 4022 initial respondents were selected from the subset of the Japanese population who were born before January 1989 (males and females aged 20 years or above as of 31st January 2009), which represented 81.5% of the total Japanese population at the time (based on statistics from February 2009). Respondents were selected via two-stage stratified sampling from the National Residents Register or the Electoral Roll for specific survey areas, with at most one individual selected from each household[61].

Repondents had the option of either completing the survey online or in person. Two types of questionnaires were given to respondents, depending on their marital status. Married respondents were given the same questions as unmarried respondents, with additional questions related to their spouse. Each year, around 85-90% of households who completed the previous questionnaire(s) continued to participate in the survey, and on average each house-

[61] For further details about the JHPS methodology and implementation, see http://www.pdrc.keio.ac.jp/en/open/about-panel.html.

hold appears 4.55 times between 2009-2013 (Table 2.2 summarises the number of observations per household and number of households per year). The JHPS is only distributed to individuals from the initial sample, so no new respondents were introduced in later years[62].

Table 2.2: Summary of JHPS

No. of obs.	Freq.	Percent
1	548	13.6
2	315	7.8
3	333	8.3
4	265	6.6
5	2561	63.7
Total	4022	100

Year	Freq.	Percent
2009	4022	25.1
2010	3470	21.6
2011	3160	19.7
2012	2821	17.6
2013	2581	16.1
Total	16054	100

2.3.1 Happiness measures

The JHPS contains hundreds of detailed questions on various topics, including education, family life, and housing. This study focuses on a subset of these survey items related to household expenditure, happiness, and household characteristics.

Starting from 2011, the JHPS included questions that asked respondents to rate their happiness over different time periods (this week, this year, their entire life) on a scale of 0 (no happiness at all) to 10 (complete happiness)[63]. The same question is asked to their spouses (if applicable). Figure 2.5 shows histograms of the respondent's and spouse's happiness for all available observations, along with the mean. The distributions are very similar across time frames for both respondents and their spouse, with a mean and mode of 6 and most ratings on the higher end of the scale.

The survey also asked for happiness over specific domains (household income, own employ-

[62]One exception is when respondents died or went missing, in which case the spouse was given the survey. The spouse's responses were documented under a new respondent ID. This situation only occurred 6 times over the 5 available survey years. The full list of new and old IDs is in the data documentation given to researchers as part of the JHPS data. These individuals were not included in the analysis of this chapter.

[63]The actual question asked was: Please provide answers as to how your feeling of happiness was during the following periods, on a scale of 0 to 10, with 0 being "having no feeling of happiness at all," and 10 being "having a feeling of complete happiness."

ment, housing, amount of leisure time, the way that leisure time is spent, and health[64]). The respondents' ratings over those domains are shown in Figure 2.6[65]. There are some differences in the distributions across domains. Most respondents seem happy with their housing and health, with a large proportion of ratings in the 8-10 range. Conversely, respondents seem more unhappy with their household income, which had the lowest mean among all domains and a larger proportion of low ratings (4 or below).

Both Figures 2.5 and 2.6 show a large density of happiness ratings at 5 (around one-quarter of ratings for any given happiness measure), which could indicate measurement issues. Respondents who are unsure of their happiness over the time period or domain, or have imperfect memory of their past responses may simply select 5 (the middle number on the response scale), so this response may not always be informative about the preference ordering between time periods. There are ways to address this bias in reporting, which are beyond the scope of this chapter. The next subsection discusses other measurement issues with happiness ratings.

[64]The exact question asked was: Please provide answers as to how you feel about the present situation regarding the following, on a scale of 0 to 10, with 0 "not at all satisfied," 5 is "neither satisfied nor dissatisfied," and 10 is "fully satisfied" (circle one).

In the questionnaire, life overall was listed as a domain even though it could encompass all the other domains. It was considered as a specific domain in this chapter and used to construct consumption happiness values in the robustness checks (see Appendix 2B). The imputed happiness values and revealed preference results using the entropy-imputed estimates are not sensitive to whether 'life overall' is considered a specific domain or not.

[65]Spouses were also asked to rate domain-specific happiness, though almost all chose not to complete or only partially complete this part of the survey. For this reason, this chapter will focus on respondents' domain-specific ratings, and will use spouse's overall ratings in the collective model analysis.

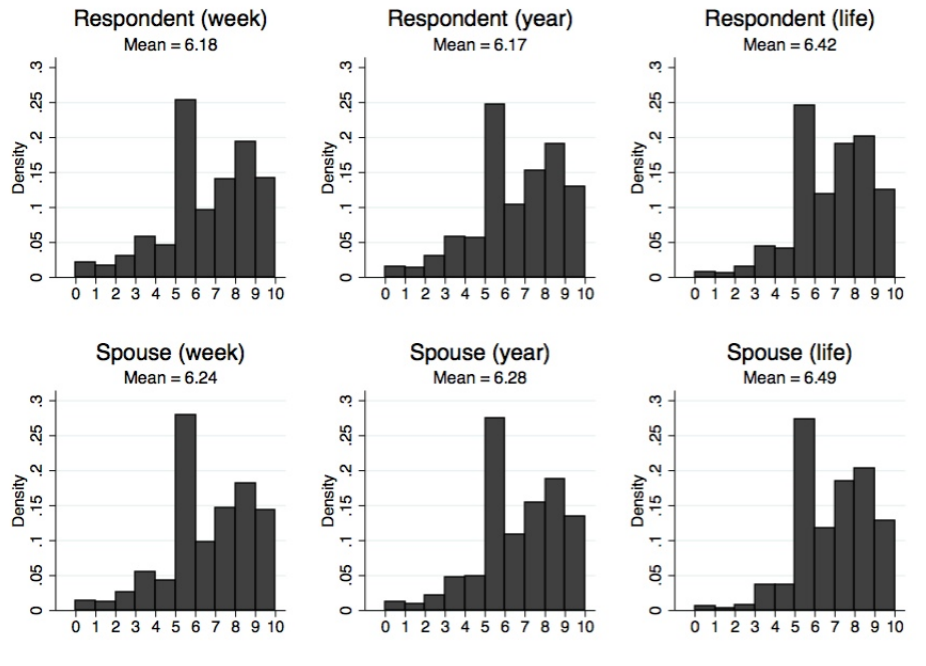

Figure 2.5: Overall happiness measures: Respondent and spouse

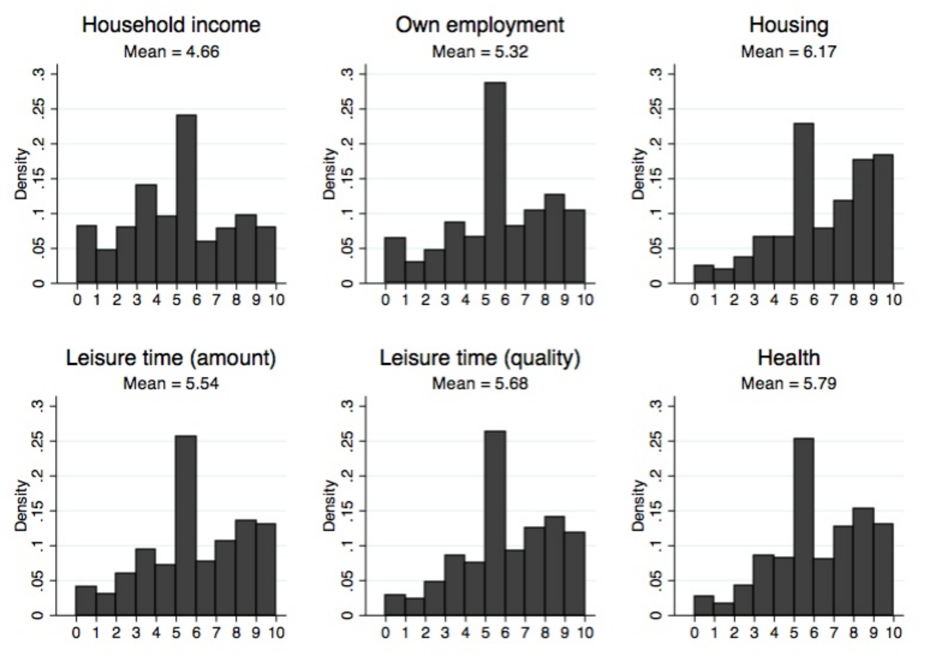

Figure 2.6: Domain-specific happiness measures: Respondent only

Table 2.3 shows that an individual's happiness measures are strongly correlated across time periods, with the highest correlation between happiness this week and happiness this year

(0.82-0.85). There is moderate correlation between a respondent's and their spouse's happiness ratings, which is expected since they are likely to experience the same household-level events. The correlation is not perfect because each person may react differently to the same event, and may also experience individual-specific events. Respondent's happiness ratings are moderately correlated across domains, with the strongest correlation being between the amount and quality of leisure time, and other correlations ranging from 0.32-0.49 (Table 2.4). Respondents' overall happiness ratings are also moderately correlated with domain-specific happiness ratings (Table 2.5).

Table 2.3: Correlations between happiness measures over different time frames for respondents (R) and their spouse (S)

Time frame	Week (R)	Year (R)	Life (R)	Week (S)	Year (S)	Life (S)
Week (R)	1					
Year (R)	**0.82**	1				
Life (R)	**0.68**	**0.76**	1			
Week (S)	0.44	0.42	0.38	1		
Year (S)	0.43	0.47	0.42	**0.85**	1	
Life (S)	0.38	0.41	0.45	**0.72**	**0.79**	1

Table 2.4: Correlations between domain-specific happiness measures for respondents

Domains	Income	Empl.	Housing	L. time	L. quality	Health
Income	1					
Empl.	0.49	1				
Housing	0.45	0.42	1			
L. time	0.37	0.35	0.44	1		
L. quality	0.37	0.41	0.45	**0.71**	1	
Health	0.32	0.40	0.38	0.37	0.48	1

Table 2.5: Correlation between respondent's overall and domain-specific measures

	Period		
Domain	Week	Year	Life
Household income	0.34	0.36	0.34
Own employment	0.42	0.46	0.40
Housing	0.34	0.36	0.36
Leisure time (amount)	0.35	0.33	0.29
Leisure time (quality)	0.40	0.38	0.36
Health	0.37	0.38	0.37

Measurement issues

Ideally, the happiness ratings $\{H_t\}_{t=1,...,T}$ would represent a household's happiness over the specific domain of consumption. However, as explained earlier, happiness ratings are only available for at most two household members (the respondent and his or her spouse). These ratings may not be an accurate measure of household happiness for larger households (60% of households surveyed had 3 or more members throughout the sample period). To explore the severity of this issue, results for 1-2 person households can be compared with those of larger households, and GARP-H was also repeated using the average or the sum of the respondent's and spouse's reported happiness as a robustness check (see Appendix 2B)[66].

Another issue is a domain mismatch, since the main happiness rating used is overall happiness rather than consumption-specific happiness. A better comparison between happiness and utility would use respondent's happiness with their consumption or living standards, but the JHPS did not include these domains in their survey. The closest domain, income, will be used for the GARP-H tests in this chapter, with additional robustness checks using imputed consumption-specific happiness shown in Appendix 2B.

2.3.2 Other variables

Summary statistics of the household characteristics used in this chapter are given in Table 2.6, with additional variables given in Tables 2.19-2.21 (Appendix 2A). These are taken from the 2009 survey year and individual-specific variables refer to the respondent (unless otherwise stated). In the survey, the term 'household head' is explicitly defined as the main earner. Most respondents were either the household head or are married to the household

[66]Note that happiness reports can only be comparable across individuals if they are a cardinal measure. Another assumption required for interpersonal comparability is that the meaning of a 1-unit increase in happiness is the same across individuals, so that comparing unweighted sums and simple averages is indicative of actual changes in household happiness across periods. These assumptions are quite strong.

head, and therefore are likely to be in charge of household decisions such as expenditures, which are the focus of this chapter. Most respondents are married, own a home, have a service-related occupation, and obtained at least a high school education. Section 2.3.4 will discuss the relationship between these household characteristics and happiness measures.

Table 2.6: Household characteristics

Variable	N	Mean	SD	Min	Max
Age (respondent)	4022	50.29	16.34	20	93
Age (household head)	3957	55.00	14.34	20	103
Household size[67]	4022	3.28	1.47	1	10
Own a home	4022	0.80	0.40	0	1
Household after-tax income[68]	3266	507.02	374.72	0	12000
Marital status[69]	4022	0.72	0.45	0	1

In the questionnaire section about household finances, respondents were asked to report their total expenditure in January on 14 categories of goods, to the nearest thousand yen[70]. This expenditure includes items bought on credit or installments, or using electronic payments. Average expenditure shares during the survey month (January) are summarised in Table 2.7. Since expenditure was only reported for one month instead of annually or over an entire year, there is a huge variation in the proportion of total household expenditure allocated to each category that might be smoothed out in annual data. For example, a large proportion of expenditure on education could be due to paying termly tuition fees due in January, or spending on medical care could be due to unforeseen health complications that only happened in January that year. Since household behaviour is not observed throughout the year, models of static utility maximisation (as outlined in Section 2.2, where households make decisions from period-to-period rather than make plans for future periods) are more

[67] Refers to the number of people (including respondent) currently living in the same house as the respondent.

[68] In 10,000's of yen i.e. 507.02 = ¥5,070,200 (approximately 48,500 USD). Household after-tax income would be reported as 0 for unemployed respondents or retired couples living alone.

[69] A dummy variable (=1 if respondent is married).

[70] Respondents were also asked to report other expenses such as pocket money and remittances, which are not included in this analysis. See http://www.pdrc.keio.ac.jp/en/open/new-about-tougou.html for full category names.

appropriate than models of intertemporal utility maximisation.

Table 2.7: Expenditure shares: Summary statistics

Variable	Mean	SD	Min	Max
Food	0.28	0.14	0	1
Mealsout	0.06	0.06	0	0.82
Fuel	0.13	0.07	0	1
Durables	0.03	0.06	0	0.94
Digital appliances	0.03	0.08	0	0.87
Clothes	0.06	0.06	0	0.91
Medical care	0.05	0.07	0	1
Transportation	0.07	0.05	0	0.92
Communication	0.07	0.05	0	1
Internet	0.02	0.03	0	0.91
Education	0.04	0.09	0	0.90
Recreation	0.06	0.08	0	0.87
Housing	0.09	0.15	0	0.99
Other	0.08	0.08	0	0.94

Consumer Price Index (CPI) data was taken from the Japanese Statistics Bureau website[71]. Both annual and monthly CPI data are available for both aggregated categories of goods (for example, recreational durable goods), as well as subcategories of items (for example, TV sets), and the aggregated CPI is calculated from a weighted average of its disaggregated components. Monthly CPI data was used for consistency with the expenditure reporting period. 2010 is the base year, so prices for each item average to 100 across all months in 2010.

All 14 JHPS expenditure categories were listed in some form in the CPI data. 9 out of these 14 expenditure categories exactly matched an aggregate category in the CPI data, while the price indices for the other 5 categories were obtained by a weighted average of the prices indices of appropriate CPI subcategories (using unscaled expenditure weights). Table 2.22 (Appendix 2A) outlines how expenditure categories in the JHPS and CPI were matched.

[71]See http://www.stat.go.jp/english/data/cpi/ for information about how the CPI was constructed, and the full list of goods included.

2.3.3 Classifying goods for the collective model

The 14 aggregate goods in the JHPS data were classified as either public or private (Table 2.8). This prior specification is required in order for Theorem 2.1 to hold (Cherchye et al 2011b). While the classification was immediately apparent for most goods (for example, clothes can only be worn by the intended household member), some categories had both a public and private element. For example, reading and recreational goods could be either private (a special interest book that only one member reads) or public (a video game or ping pong table that the whole family can play together). Following the approach of Cherchye et al (2011b), goods that could be either private or public were classified as public[72].

Table 2.8: Classification of JHPS categories

Public goods	**Private goods**
Fuel	Food
Durables	Meals out
Digitals	Clothes
Transportation	Medical
Internet	Communication
Reading & Rec	Education
Housing	Other

There are two assignable private goods for the JHPS: clothes (which includes footwear) and education. Aside from reporting aggregate expenditure on these goods, respondents are asked to break down this amount into the share (percentage) of total expenditure that came from 'common expenses', or was contributed by themselves, their spouse, children, and other household members (if applicable)[73]. Since happiness data is only available for respondents and their spouse, the revealed preference tests for the collective model will be conducted on

[72]These assumptions on public vs. private goods can be tested with 4 or more observations per household. Cherchye et al (2013) show that public goods and private goods have different testable implications even if researchers only observe the aggregate group consumption, but these differences only arise for 4 or more time periods. Since happiness data was only available for 3 years, it is not possible to test the assumptions on goods classification here.

[73]Due to the common expenses option in the questionnaire, the reported own/spouse expenditure shares provide a lower bound on the assignable quantity rather than an equality.

the subsample of two-person households.

Of the 878 two-person households remaining in 2011, 619 households had happiness responses for both respondent and spouse (20% of all households remaining in 2011). Assignable quantity information was reported by 360 households for clothing only[74], and 14 households reported information for both clothing and education.

2.3.4 Happiness and economic variables

As mentioned in Section 2.1, previous studies have estimated 'happiness functions' by looking for econometric relationships between happiness ratings and economic variables such as income and employment status[75]. Table 2.9 shows some 'happiness functions' estimated using pooled OLS, for both household heads and their spouse (where applicable). The following household characteristics were used: household after-tax income in millions of yen (net income), age of the household head (age), whether the household owned a home (homeowner), whether the person was married (married), household size (HH size), and total monthly expenditure in millions of yen (total exp). Overall happiness ratings for the year were used[76]. For both household head and their spouse, household after-tax income has a significant positive but non-linear relationship with happiness, though the effects may not be large in absolute terms[77]. Married respondents are also significantly happier than single respondents.

[74] None of the households reported information for education only.

[75] See, for example, Ferrer-i-Carbonell and Frijters (2004).

[76] Results for happiness measures over other time frames are qualitatively similar so are not reported here. Total expenditure was also included as an explanatory variable but was insignficant in all cases so was omitted from Table 2.9 and 2.10.

[77] Holding other things constant, a one-point increase in happiness would require a ¥10-13 million (96,000-124,000 USD) increase in annual after-tax income.

Table 2.9: 'Happiness functions' (clustered standard errors in brackets)

	Head	Spouse
Net income	0.101	0.076
	(0.016)	(0.017)
Net income^2	-0.001	-0.001
	(0.000)	(0.000)
Age	0.004	-0.009
	(0.003)	(0.003)
Homeowner	0.157	0.195
	(0.099)	(0.113)
Married	0.577	
	(0.118)	
HH size	-0.055	-0.066
	(0.032)	(0.036)
Year	-0.013	-0.070
	(0.026)	(0.027)
Constant	5.066	6.871
	(0.210)	(0.254)
R^2	0.033	0.011

Table 2.10 shows that there is a statistically significant relationship between domain-specific happiness and outcomes related to that domain. Household income and total monthly expenditure both have larger positive effects on domain-specific happiness than on overall happiness[78]. Similarly, home ownership is associated with a 1.6-point increase in a respondent's happiness with housing, compared to a 0.2-point increase in overall happiness. However, having to share household income or housing with more members has a significant negative effect on domain-specific happiness. These results suggest that the relationship between happiness measures and economic outcomes is worth exploring further.

[78]The positive effect of a ¥1-million increase would be 3-5 times larger on domain-specific happiness than overall happiness. The small but statistically significant coefficients on the squared variables suggest there is a non-linear relationship between income or expenditure and domain-specific happiness.

Table 2.10: Domain-specific regressions (clustered standard errors in brackets)

	Domain		
	HH income	Own Employment	Housing
Net income	0.352	0.178	0.095
	(0.023)	(0.018)	(0.018)
Net income^2	-0.004	-0.002	-0.001
	(0.001)	(0.000)	(0.000)
Age	0.015	0.007	-0.005
	(0.003)	(0.003)	(0.003)
Homeowner	0.350	0.213	1.593
	(0.116)	(0.116)	(0.114)
Married	0.274	0.466	0.007
	(0.102)	(0.108)	(0.100)
HH size	-0.298	-0.164	-0.204
	(0.032)	(0.033)	(0.033)
Year	-0.006	-0.048	-0.061
	(0.026)	(0.029)	(0.025)
Constant	2.596	4.285	5.547
	(0.229)	(0.236)	(0.222)
R^2	0.134	0.043	0.077

2.4 Results

Table 2.11 summarises all the revealed preference tests conducted in this section, and their abbreviations.

Table 2.11: Summary of revealed preference tests

Model	Happiness measure used	Abbreviation
Unitary	None	GARP
	Respondent's happiness (overall or income)	GARP-H
Collective	None	GARP-C
	None, with assignable quantities	GARP-C1, GARP-C2
	Respondent and spouse's happiness	GARP-CH
	Respondent and spouse's happiness, with assignable quantities	GARP-CH1, GARP-CH2

2.4.1 Unitary household behaviour

As discussed in the Introduction, if household behaviour is consistent with utility maximisation under the unitary model, then there must be a solution to the set of linear inequalities given by Afriat's Theorem. The price and quantity data required for the revealed preference tests were obtained from reported household expenditures and monthly CPI data. Relative prices for each year were calculated by dividing the price in year t by the price in the base year (2009 in this case)[79]. Quantities were then calculated by dividing category-specific expenditure by these category-specific prices i.e. $q_t^k = \frac{x_t^k}{p_t^k}$. These constructed prices and quantities were then substituted into the linear inequalities for the revealed preference test.

The revealed preference test gives a binary result, where 1 corresponds to 'Pass' (behaviour is consistent with utility maximisation according the given model) and 0 corresponds to 'Fail' (behaviour is not consistent with the model). Table 2.12 reports the pass rates of the GARP and GARP-H tests (standard errors in brackets), which are defined as the proportion of households whose behaviour was consistent with the model. For comparability between GARP and GARP-H, only data from 2011-2013 (up to 3 periods per household) was used for the GARP test. Among the reported measures of happiness, the respondent's overall happiness for the year (rather than week or entire life) was used since it corresponded most closely with the time frame of the expenditure data[80], as well as domain-specific happiness over income. Almost all households (98.7%) behaved consistently with the unitary model, but most household behaviour was not consistent with their preference ordering given by either set of happiness numbers.

It is possible that the respondent's happiness is not representative of the entire household's happiness. Table 2.12 also presents the pass rates for single-person households[81]. The proportion of single-person households that pass any of the happiness tests is slightly higher but not significantly different from the proportion of all households that pass, so interpreting happiness numbers as a preference ordering over consumption appears to be problematic

[79] Thus the price of good k in year t was calculated as $p_t^k = \frac{CPI_t^k}{CPI_{2009}^k}$.

[80] As a robustness check, GARP-H was repeated using overall happiness measures from different time frames. The results are very similar and so are not reported here, but are available upon request.

[81] Since household size can change over the survey period, and changes or anticipated changes in household size may affect reported happiness in current and subsequent periods, these results are for the subset of households that consisted of a single person throughout the full survey period (2009-2013).

even for individuals.

As discussed in Section 2.3.1, happiness reports are available for at most two people in each household. As a robustness check, the GARP test was repeated using the average and sum of the respondent's and spouse's report (see Table 2.25, Appendix 2B), but gave similar pass rates as GARP-H[82].

Table 2.12: Pass rate comparison (2011-2013, standard errors in brackets)

	All households	1-person
GARP	0.987	0.994
	(0.113)	(0.077)
GARP-H (overall)	0.086	0.101
	(0.280)	(0.303)
GARP-H (income)	0.090	0.114
	(0.286)	(0.318)

Table 2.13 categorises households according to their GARP and GARP-H results. The majority of households (83.7%) passed GARP, but most failed GARP-H for both happiness measures. Fewer than one-fifth (15.0%) of households behaved consistently with the preference ordering given by their happiness numbers (either overall or domain-specific). These test results provide strong evidence of inconsistency between preferences implied by happiness measures and utility maximisation.

Table 2.13: Unitary model test results (all households)

	Fail GARP-H (overall)		Pass GARP-H (overall)	
	Fail GARP-H (inc)	Pass GARP-H (inc)	Fail GARP-H (inc)	Pass GARP-H (inc)
Fail GARP	1.3%			
Pass GARP	83.7%	6.4%	6.0%	2.6%

[82]Measures of household happiness constructed using both the respondent's and spouse's report might be more informative of actual household happiness for 2-person households compared to larger households if happiness measures can be compared across individuals (1 and 2-person households account for 34% of all households in 2009). However, the last column of Table 2.25 (Appendix 2B) shows that using this additional information did not significantly improve the pass rate of 2-person households, so it is likely that the low pass rates are largely due to the preferences implied by the happiness numbers rather than the particular measure of happiness used.

Evaluating pass rates

Despite the lower pass rate, could a model that restricts preferences according to happiness ratings still be 'better' than the standard model? The pass rate alone cannot indicate whether one model is better at explaining observed behaviour compared to another. A model with a higher pass rate may simply be less restrictive: it would be easy to get a 100% pass rate with a model that allowed any sort of behaviour. The model must therefore put at least some observable restrictions on the data so that it would be possible to make choices that are consistent as well as inconsistent with this structure. That is, for a 'pass' result to be taken seriously, a 'fail' result must also have been possible.

One concern with the standard unitary model is that for some sets of budget constraints, it is impossible for the consumer to choose in a way that violates utility maximisation. Figure 2.7 (left panel) and Figure 2.1 from Section 2.2.1 illustrate some examples. In these examples, the two budget constraints do not cross, so any choice the consumer could make will satisfy GARP and hence be consistent with utility maximisation[83]. Happiness measures add more structure to the data, so that it is possible to find a violation even in this situation (Figure 2.1, right panel).

The pass rates therefore should be evaluated in light of how likely it is to find an inconsistency with the model for the given expenditure constraints. If relative prices (prices across categories) do not vary much over time or are increasing proportionately over time (for example due to inflation), then budget constraints will rarely or may never cross, so finding a rejection of GARP would be virtually impossible. Power is one measure of the likelihood of making an inconsistent choice. Measuring the power of each household's test would give an indication of how prevalent such issues are in the data.

[83] Recall that GARP means that it is possible to find a rational preference ordering between the consumption bundles chosen. In this case, any choice on line A will always be strictly preferred to any choice on line B. Given these budget constraints, it is not possible to make contradictory choices i.e. choose a point on A that is strictly preferred to a point on B but vice versa.

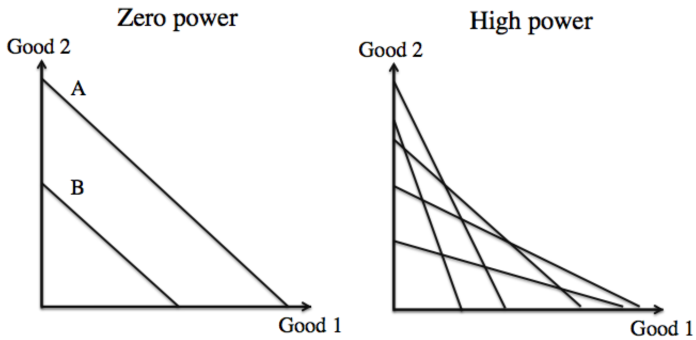

Figure 2.7: Power examples

The power of a test is defined as the probability of rejecting a false null hypothesis i.e. Pr(reject $H_0|H_0$ is false). The null hypothesis in this case is that the behaviour is consistent with the model. There is no obvious alternative hypothesis (if consumers are not utility maximising then it is unclear how they would behave instead), but the common alternative is uniform random behaviour[84]. The left panel of Figure 2.7 is an example of a test with zero power: it is impossible to reject the null hypothesis, since any choice is consistent with GARP. The right panel shows an example of high power: among all possible choices there are many that violate GARP and only a few that satisfy GARP.

Power was calculated separately for each household using the method of Bronars (1987)[85]. Table 2.14 compares the power and pass rate (standard errors in brackets)[86]. Based on pass rate alone, GARP appears to be better than GARP-H at explaining the observed behaviour. However, when considering the power of the tests, it is not clear which model is better. GARP is far less restrictive than GARP-H, with the average chance of finding inconsistent behaviour

[84]This alternative was first proposed by Becker (1962), and assumes that the consumer randomly selects one point on each budget constraint.

[85]One point was chosen from each budget constraint by generating expenditure shares that are uniformly distributed, and the revealed preference test was run using the prices (constructed as before) and the quantities constructed from the sampled budget shares. This process was repeated 1000 times per household, with the test result ('pass' or 'fail') recorded each time. Estimated power is the proportion of times in which the generated data failed the test.

[86]Since power was only measured to 3 decimal places, an estimate of 0 does not necessarily mean that it is actually impossible to reject the null hypothesis, but rather that it is very difficult to find a rejection. Likewise, a power estimate of 1 means that it is very difficult to find data that passes the restrictions given by the model.

at 2% compared to 91% for either happiness measure. There appears to be a tradeoff between pass rate and power: the standard unitary model (GARP) has a high pass rate but generally low power, whereas the added restriction on happiness numbers gives GARP-H low pass rates but much higher power on average.

Table 2.14: Pass rate, power, and predictive success (all households, 2011-2013)

	Pass rate	Power	PS
GARP	0.987	0.022	0.009
	(0.113)	(0.052)	(0.115)
GARP-H (overall)	0.086	0.842	-0.066
	(0.280)	(0.151)	(0.271)
GARP-H (income)	0.090	0.842	-0.072
	(0.286)	(0.151)	(0.264)

Intuitively, a 'better' model should have less of a tradeoff between pass rate and power. A model that is good at explaining behaviour would put many restrictions on the choices that consumers can make (high power), and observed choices would satisfy these restrictions (high pass rate). Conversely, a model that is very bad at explaining behaviour would have a very large tradeoff between pass rate and power: it would either make very general predictions about how consumers behave that are difficult to falsify (high pass rate and low power), or would put restrictions on behaviour that are not observed in the data (low pass rate and high power).

A formal way to evaluate this tradeoff is to use a measure of predictive success (PS). This measure uses both pass rate and power to compare the explanatory ability of models. There are many measures of predictive success[87], but the common measure used is PS = Pass rate - (1 - power), where higher values indicate models that are better at explaining the data.

[87] See Selten (1991) and Beatty and Crawford (2011) for an overview and discussion. Formally, measures of predictive success should satisfy the following properties: 1. Monotonocity (Data that is consistent with a restrictive model should get a higher value than data that is not consistent with an unrestrictive model), 2. Equivalence (the measure should assign the same value when data is consistent with a model that can rationalise any behaviour, as the situation when data fails a model that cannot explain any observed behaviour), and 3. Aggregability (the group-level measure equals a weighted sum of individual-level measures).

(1 - power) can be interpreted as the range of behaviour that the model allows: if power is zero, then the model allows any sort of behaviour, but if the power is close to 1, then the model allows only some specific patterns of behaviour.

PS ranges between -1 and 1: $PS \to 1$ when the data satisfies the conditions of a very restrictive model, while $PS \to -1$ when the data fails to satisfy the conditions of a model that allows almost any observed behaviour. $PS = 0$ when the pass rate increases by the same amount as the power decreases, which is what would happen if the restrictions imposed did not add any explanatory ability to the model.

The final column of Table 2.14 gives summary statistics for PS for the three models. On average, there is no clear benefit from including happiness numbers in revealed preference tests: the pass rate decreases more than proportionately than the increase in power. PS becomes negative when happiness restrictions are imposed, but PS is very close to 0 for all three models. Therefore, models that include happiness measures are worse at explaining observed choices, but not significantly worse. Also, all models are generally not good at explaining the data: the low standard deviation indicates that PS is close to 0 for most households for all models.

Optimisation error

The revealed preference test result is binary (either 'pass' or 'fail'), so even if a household's choices were almost consistent with the model they still received a 'fail' result. In reality, survey data might contain measurement error (especially if expenditure reports are based on memory or if reports are rounded to the nearest ten/thousand), or households may make 'human errors' (optimisation errors). To evaluate a model's ability to explain behaviour, it therefore matters how 'badly' the data fails the test. If a slightly different set of choices or expenditures would have been consistent with the model (a 'close' failure), then it is plausible that optimisation errors are the source of failure. However, if even a wildly different set of expenditures is still not consistent with the model (a 'bad' failure), then it is plausible that the wrong model is being used.

Optimisation errors can be modeled as inefficient spending. The household is not choosing optimally because it is not spending as much as it can (for whatever reason). The Afriat efficiency index[88] (e) is a measure of how 'close' the data are to passing a test, and is defined as the minimal perturbation (expressed as a proportion) of every budget constraint that allows the data to pass the test[89]. $e \in [0, 1]$ can be interpreted as the overall efficiency of a household's choice behaviour: $e = 1$ for households that pass the standard revealed preference test, $0 < e < 1$ means that the household passes the test if all of their budget constraints are multiplied by the factor e, and $e = 0$ is a situation where any observed behaviour is rationalisable (households are allowed to be completely inefficient in their spending).

$(1 - e)$ can be interpreted as the extent of optimisation error that is required for the data to pass the test. If the required optimisation error is quite large (e is very small), it is reasonable to conclude that the observed behaviour is not consistent with the model. If e is very large (households were close to being efficient with their spending) then the observed behaviour could be consistent with the model. Figure 2.8 shows the intuition behind the Afriat efficiency measure. Here, allowing the household to be slightly inefficient in one period (an inward parallel shift of the steeper budget constraint to the grey line) would make their behaviour consistent with utility maximisation.

[88]Devised by Afriat (1967, 1972) and Varian (1985, 1990).

[89]Formally, if the data passes GARP at a particular efficiency level e_1, then it is possible to construct a preference relation $R(e_1)$ such that if $\mathbf{q}_t R(e_1) \mathbf{q}_s$, then $e\mathbf{p}_t\mathbf{q}_t R(e) \mathbf{p}_s\mathbf{q}_s$. The Afriat efficiency index $e \in [0, 1]$ is the largest number that satisfies this condition.

Figure 2.8: Finding Afriat efficiency (e)

The Afriat efficiency index was calculated (to the nearest 0.05) for each household for the three revealed preference tests (GARP, GARP-H (overall), and GARP-H (income)). For households that passed the revealed preference test, $e = 1$. For households that failed the revealed preference test, one budget constraint was multiplied by a proportion (starting with 0.95 and decreasing in increments of 0.05), holding the happiness numbers constant, and GARP-H was repeated until the data passed. This process was repeated for each budget constraint, and e was taken to be the maximum across all budget constraints[90].

Figure 2.9 shows the distribution of e for GARP and GARP-H (income)[91]. All households that failed GARP only required a small movement of one of their budget constraints to pass the test[92], so optimisation error is a plausible cause of failure. In contrast, only 23% of households would pass GARP-H (income) allowing for reasonable optimisation errors of 5-10%, with most households that failed GARP-H (income) having very low cost efficiency, suggesting that optimisation error in expenditure is unlikely to be the reason for the low

[90] e in this case is the lower bound of the actual Afriat efficiency number; it is possible that the data are rationalisable using a number in the interval $(e, e + 0.05)$. However, intervals of 0.05 are sufficiently small for the purposes of this descriptive exercise.

[91] The distribution of GARP-H (overall) is very similar to that of GARP-H (income) so is not reported here.

[92] For GARP, cost efficiency ranged from 95-100%, with $e = 0.95$ for all households that failed GARP. For a household with the mean monthly total expenditure in 2009 of ¥250,000 (2400 USD), $e = 0.95$ implies a 'wastage' of ¥13,000 (125 USD), which seems plausible.

pass rates.

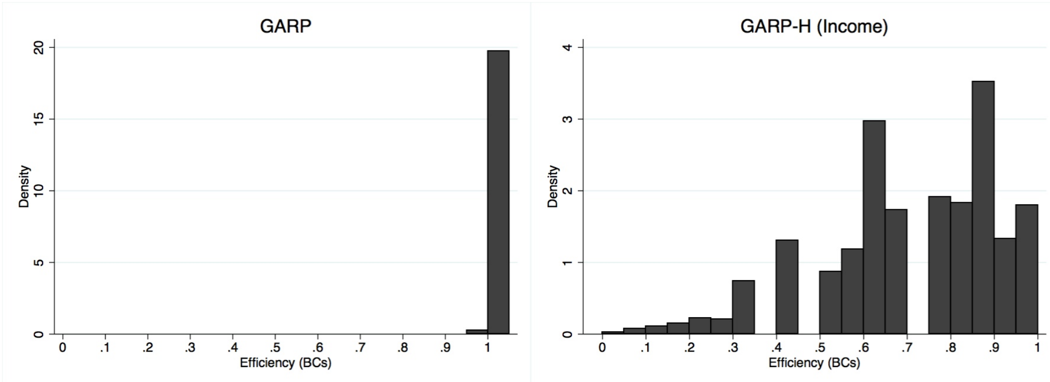

Figure 2.9: Afriat efficiency (e): Optimisation error

Reporting error

It is also possible that the inconsistency between the two preference orderings is due to misreporting of the happiness numbers. The reported happiness numbers may not represent a household's true preferences for various reasons: households may not have remembered exactly what number they reported last year[93], or simply circled the wrong number on the questionnaire. Reporting errors could be the reason why the Afriat efficiency is very small or zero for most households, because misreporting would cause reported preference orderings to differ from actual preferences (and observed behaviour) and may require large inefficiencies in spending to rationalise.

The possibility of misreporting can be investigated using a minimum distance measure d_{min}, which is taken to be the smallest sum of squared deviations from data-consistent happiness numbers, out of all possible consistent combinations[94]. This minimum distance mea-

[93] Although the happiness numbers do not have a cardinal meaning, to express relative preferences over periods, households must remember what they reported last year in order to report a number that correctly represents their preferences. For example, suppose an individual prefers period 1 to period 2 and wanted to express this preference using numbers. If he reported a '6' in period 1 but misremembered reporting a '7' last year, then reporting a '6' in period 2 would be a correct representation of his preferences from his perspective, but would be inconsistent with the actual data.

[94] For example, if $\mathbf{H} = (8, 5, 7)$, but period 2 $\succ$ period 1 $\succ$ period 3 according to consumption data, the closest data-consistent happiness numbers would be $(8, 9, 7)$ and $d_{min} = (8-8)^2 + (9-5)^2 + (7-7)^2 = 4$.

sure is easier to interpret than an Afriat efficiency measure, because happiness ratings use a discrete scale so reporting errors would be integers (for example, reporting 5 instead of 6) rather than proportions of the original ratings[95]. Figure 2.10 shows the actual distribution of d_{min} using income-specific happiness, along with a simulated distribution based on randomly-generated happiness numbers[96]. $d_{min} = 0$ for households that pass GARP-H (income), and for three time periods, $max(d_{min}) = 3(11^2) = 363$. Apart from a few outliers, the degree of reporting error is quite small, ranging from 0-10 for most households, which roughly corresponds to misreporting a happiness number by 1 or 2 points each time (for example, reporting a 6 instead of a 7). According to the simulated distribution, if the happiness numbers were completely uninformative about preferences, the minimum difference would be in the 15-25 range for most households (misreporting of 2-3 points each time). These results indicate that income-specific happiness numbers are somewhat informative about preferences over consumption, and that reporting error is a plausible explanation for the low pass rates of GARP-H.

[95] Also, the minimum distance measure treats households the same regardless of their average happiness levels, meaning that a 1-unit error would give the same value of d_{min} regardless of whether the household was very happy (reporting a 9 instead of a 10) or very unhappy (reporting a 1 instead of a 2). This characteristic matters more for happiness compared to expenditure, because there are more values in the extreme ends of the distribution (see Figure 2.6).

[96] To generate the simulated distribution, numbers from 0-10 were randomly drawn with replacement and d_{min} was calculated. For each household, this process was repeated 30 times and the average of d_{min} was taken. The simulated distribution shows the means of d_{min}. Again, the results are similar when using overall happiness instead of income-specific happiness, and so are not reported here.

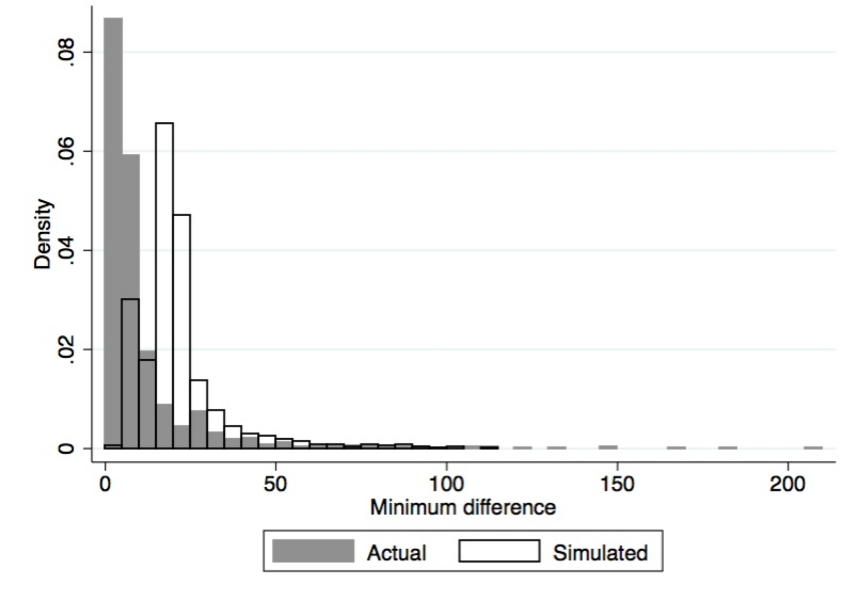

Figure 2.10: Minimum distance: Actual vs. simulated distribution

Under the unitary approach, the preference orderings implied by happiness numbers are generally not consistent with revealed preferences. This result is robust to the happiness measure used (both the time frame asked in the survey and constructed measures). As shown through power and predictive success measures, adding information about happiness to the standard unitary model does not improve its ability to explain observed behaviour.

However, these findings could be due to the assumption of unitary household behaviour, since individual preferences might plausibly differ within a household and depend on personal rather than household-level consumption. The next section uses the collective model of household behaviour to determine whether or not the results from this section still hold in settings where each household member's preferences are allowed to vary and depend on personal consumption.

2.4.2 Collective household behaviour

The analysis in this section is restricted to households that had exactly two members throughout 2011-2013. Six different variations of the collective household model were tested: without happiness data or assignables (C), without happiness data and one assignable (C1), without happiness data and two assignables (C2), and their counterparts with happiness data (CH, CH1, and CH2 respectively). Since domain-specific happiness is not significantly better than overall happiness (based on pass rate, power, and predictive success), the happiness measure used was the respondent's overall happiness rating. Due to many missing domain-specific observations for spouses, there were no households that had complete domain-specific responses from both respondent and spouse for more than 1 year, making GARP-H (income) infeasible for the collective model.

The revealed preference test was done computationally using the mixed-integer program described in Section 2.2.2. As shown in Table 2.15, the collective model is consistent with a greater proportion of 2-person household behaviour compared to the unitary model (100% compared to 98.7% without using happiness numbers, and 37.8% compared to 9.1% when using happiness numbers). However, using happiness numbers to fully determine each household members' preference ordering between periods dramatically reduces the pass rate: only 37.8% of all data that could be rationalised without assignable information was consistent with the happiness ordering, with similar reductions for the variations that use assignable information.

Table 2.15 summarises the pass rate, power, and predictive success of these six variations of the collective household model, with the results of the unitary approach from Section 2.4.1 presented for comparison (standard errors in brackets). As discussed in Section 2.2.2, the mixed-integer program for the collective model is much more computationally intensive than the linear program for the unitary model[97], so power was estimated as the value that the pass rate converged to after at least 100 random draws from the data[98].

[97] Testing for C, C1, and C2 for a 2-person household with 3 observations involves checking over 500 possible preference orderings for each random draw from a particular household.

[98] Since the power estimate is an average ($\frac{\text{\# random draws that pass}}{\text{\# of random draws}}$), after many random draws this average should converge to a consistent estimate of the actual power. The power estimate is there-

Among variations that use the same happiness information, including available information about assignable private goods also lowers the pass rate and increases the power of the test, especially in the absence of happiness numbers. This tradeoff indicates that the additional restrictions on preference orderings and/or assignable quantities cannot help explain behaviour. However, the measure of predictive success ranges from -0.1 to 0.1 across all models, so on average, all models are not very good at explaining the observed behaviour.

Table 2.15: Summary of results (2-person households)

	Pass rate	Power	PS	N
CH	0.378	0.620	-0.002	619
	(0.020)	(0.019)	(0.000)	
CH1	0.369	0.684	0.053	360
	(0.025)	(0.020)	(0.011)	
CH2	0.143	0.735	-0.122	14
	(0.097)	(0.081)	(0.084)	
C	1.000	0.007	0.007	619
	(0.000)	(0.000)	(0.000)	
C1	0.758	0.370	0.129	360
	(0.023)	(0.015)	(0.016)	
C2	0.357	0.517	-0.126	14
	(0.133)	(0.079)	(0.115)	

To summarise the findings of this chapter, Table 2.16 organises 2-person households according to their revealed preference test results, with percentages in brackets[99]. Most household behaviour is consistent with the unitary model (GARP), but not consistent with the preference ordering given by happiness numbers. The collective model was able to explain the behaviour of all households that failed both unitary models (GARP and GARP-H), but most of these households also do not behave according to the preference ordering given by their

fore considered to converge if the difference in this average between consecutive random draws is less than a pre-set tolerance level (0.001 in this case). To prevent the estimation process from continuing indefinitely, an upper limit was set for the total number of draws (1000), after which the power estimate for that household was taken to be $\frac{\text{\# random draws that pass}}{1000}$ (which would then be the same power measure used in Section 2.4.1). Another issue with this method is that power could be estimated as 0 or 1 if the first few consecutive draws all fail or pass (respectively), so a lower limit of 100 was set for the number of draws.

[99]Note that GARP-C is not included because all households passed.

happiness responses. Overall, these findings suggest that happiness measures and revealed preferences are generally not consistent.

Table 2.16: Summary of main results (2-person households)

	Fail GARP-H (inc) Fail GARP-CH	Pass GARP-CH	Pass GARP-H (inc) Fail GARP-CH	Pass GARP-CH
Fail GARP	0.8%	0.5%		
Pass GARP	55.6%	34.4%	5.8%	2.9%

2.4.3 Utility, household characteristics, and well-being

Pass rates may depend on observable variables. It is plausible that households that face stricter budget constraints are more likely to plan expenditure carefully in order to maximise their available income. To check if any household variables significantly affected the probability of passing any of the revealed preference tests, probit regressions were run on the test results and the following standard household characteristics: age of the household head (Age), educational qualifications of the household head (Educ), occupation type of the household head (Occ type), household size (HH size), dummy variable for home ownership (Homeowner), household's total after-tax income in 2008 (Net income, in millions of yen), a dummy variable for marital status (Married)[100], and the number of observations available for that household (Num). This regression only includes households in which the respondent or his/her spouse was the head of the household in 2009, since some variables such as education were only asked for the respondent and his/her spouse, and so are not available if the household head was another person.

Table 2.17 shows the coefficients from the probits on the GARP, GARP-H, and GARP-CH test results, which represent the effect of variables on the probability that behaviour is consistent with the given model[101]. None of these household characteristics are significant, ex-

[100]These variables were all from the 2009 survey.
[101]Probits were also run for the other measures of happiness. Aside from the number of observations

cept for the number of observations for that household (Num). It is sometimes harder for data to be consistent with the model when there are more observations because there are usually more observable implications that must be satisfied. There is generally greater potential to detect a violation when there are more budget lines since there is also more information about preference orderings, provided that these budget lines intersect. As with linear models, the pseudo-R^2 measures the proportion of variation in the data that the model can explain, relative to an intercept-only model[102]. The pseudo-R^2 of all probits are low, so household characteristics do not seem to predict consistency with utility maximisation (with or without happiness measures).

Table 2.17: Probit coefficients (3 d.p., standard errors in brackets)

	GARP	GARP-H	GARP-CH
Net income	-0.001	0.000	0.006
	(0.008)	(0.000)	(0.026)
Age	0.006	-0.009	0.013
	(0.007)	(0.005)	(0.008)
Homeowner	-0.188	0.068	0.045
	(0.275)	(0.130)	(0.229)
Married	-0.448	0.131	-0.285
	(0.372)	(0.172)	(0.445)
HH size	0.015	-0.065	0.079
	(0.051)	(0.044)	(0.164)
Educ	-0.034	-0.018	-0.042
	(0.056)	(0.044)	(0.074)
Occ type	-0.116	-0.045	0.028
	(0.129)	(0.072)	(0.113)
Num		-1.387	-0.786
		(0.117)	(0.323)
Constant	2.852	3.151	1.290
	(0.614)	(0.479)	(1.162)
Pseudo-R^2	0.023	0.175	0.033

Another question of interest is: Are households whose happiness ratings are consistent with utility maximisation actually happier? If utility and happiness are positively correlated, then households who act rationally may be happier than households who do not.

Table 2.18 shows coefficients from regressing the respondent's and spouse's happiness rat-

for that household (Num), none of the other variables were significant predictors of the test result.

[102] As with linear models, if the model can explain all the variation in the data, then the pseudo-R^2 would be 1.

ings used in the revealed preference tests on their respective test result (GARP-H (income)) and household characteristics[103]. Rationality is a positive predictor of overall happiness in both the unitary and collective model, but is only significant for respondent's yearly happiness in the collective model. The regression of the respondent's income-specific happiness shows that happiness over income is significantly correlated with variables besides income, and therefore is a broader measure of income-related satisfaction than the standard utility function used in Section 2.2, which only depends on consumption. These results suggest that there is not always a strong relationship between happiness and utility, so one measure cannot always be used as an indicator or proxy for the other.

Table 2.18: Rationality and happiness (clustered standard errors in brackets)

	Unitary	**Collective**	
	Respondent	Respondent	Spouse
Pass	0.209	0.359	0.256
	(0.141)	(0.154)	(0.158)
HH size	-0.200	0.524	0.071
	(0.032)	(0.154)	(0.160)
Age	0.013	0.009	0.006
	(0.003)	(0.007)	(0.006)
Homeowner	0.534	0.324	0.571
	(0.120)	(0.237)	(0.235)
Married	0.359	1.256	
	(0.106)	(0.384)	
Net income	0.361	0.141	0.138
	(0.037)	(0.046)	(0.041)
Net income^2	-0.009	-0.003	-0.002
	(0.001)	(0.001)	(0.001)
Year	0.001	-0.023	-0.063
	(0.024)	(0.047)	(0.047)
Constant	3.050	2.846	4.972
	(0.229)	(0.741)	(0.509)
R^2	0.041	0.040	0.030

2.4.4 Extensions

One limitation of this study was that the happiness questions were only recently introduced to the JHPS data, which reduced the number of observations per household that could be

[103]The regression results for GARP-H (overall) are qualitatively similar and so are not reported here.

used for the revealed preference tests. A potential concern is that it would be harder to find a rejection of GARP with fewer observations per houseshold, however this was not a big issue since the power of the tests with happiness restrictions was already quite high and pass rates were low even with only 2 observations per household. Still, with more observations per household, it would be possible to test the assumptions made in Section 2.3.3 about the public/private nature of expenditure categories, to determine how sensitive the results for the collective model are to these prior specifications.

Another issue with the JHPS data is that expenditure was only reported for one month each year, so it was only possible to test consistency for static models. The static models considered here had very high pass rates (98-100%) and low power, meaning that there were few restrictions placed on the data. It would be interesting to check if these results hold for more restrictive models; with annual expenditure data it would be possible to use intertemporal models.

While the results of this chapter were robust to the happiness measures used from the survey, none of the measures used included respondent's happiness from consumption. It is very difficult to find a survey that has both the expenditure data required for revealed preference tests and domain-specific happiness ratings for consumption or living standards, but for greater potential comparability between happiness measures and preferences, it would be ideal to repeat these revealed preference tests using happiness measures that are consumption-specific, such as happiness over material goods or possessions.

2.5 Conclusion

This chapter investigates whether or not happiness measures and revealed preferences are consistent in an economic context. This issue has important policy implications: if happiness over consumption is consistent with actual consumption patterns, then policymakers can use an individual's happiness ratings to construct his or her complete and unique preference ranking over time periods. Since revealed preference orderings are usually partial and non-unique, happiness measures could help policymakers make more accurate demand predictions and welfare analysis. Such consistency would also support the policy literature approaches that use happiness measures as representative of utility, or to construct economic measures of well-being such as social welfare.

Nonparametric revealed preference tests were conducted using individual-specific happiness measures and household-level expenditure data from the Japanese Household Panel Survey (JHPS). These revealed preference tests use observable data only to determine whether behaviour is consistent with a given model. Two common models of household behaviour were used: unitary (household members have the same preferences) and collective (household members have different preferences). To test for consistency with revealed preference, the preference orderings given by happiness numbers were included into standard revealed preference tests as additional constraints. Household behaviour that was utility-maximising and also satisfied the additional constraints would be consistent with their happiness measures.

In this economic context, happiness measures are generally not consistent with revealed preferences. Although the expenditure behaviour of most households is consistent with utility maximisation, the preference orderings implied by their happiness ratings are not consistent with utility-maximising choices. These findings are similar across different models of household behaviour and happiness measures, and are largely not due to optimisation error, suggesting that happiness and utility are two distinct and non-interchangeable measures of well-being.

From a modelling perspective, there is little justification for using happiness measures to impose additional restrictions on the data. Doing so increases the power of revealed preference tests but lowers the pass rate more than proportionately, which limits the usefulness of these orderings. When evaluated according to a measure of predictive success, models that include preference orderings given by happiness ratings are slightly worse at explaining behaviour compared to their standard counterparts.

Overall, this analysis suggests that taking happiness measures as an indicator or proxy for revealed preferences or utility is a very strong assumption in economic contexts. Happiness measures for economic variables such as leisure time and living standards can be useful supplements to standard economic measures for evaluating well-being and progress in society. However, due to their general inconsistency with revealed preferences, policymakers should be careful when using happiness measures as representative of utility, or to construct economic measures of well-being or preferences.

Appendix 2A: Supplementary tables

Table 2.19: Respondent's position in household

	Not married	Married	Total
Household head	512	1521	2033
Household head's spouse		1237	1237
Other	598	117	715
Total	1110	2875	3985

Table 2.20: Highest educational attainment of household head (as of 2009)

Education type	Freq.	Percent
Junior high school	368	11.3
High school	1,440	44.23
Junior college[104]	229	7.03
4-year university	985	30.25
Graduate school	83	2.55
Other	151	4.64
Total	3,256	100

[104] Includes specialised school.

Table 2.21: Occupation type of household head (in 2009)

Occupation type	Freq.	Percent
Primary sector[105]	78	2.95
Manufacturing	491	18.56
Services[106]	1,017	38.44
Services (requiring further education)[107]	1,060	40.06
Total	2,646	100

Table 2.22: Corresponding categories in JHPS and CPI data

JHPS category	CPI category	CPI subcategories
Food	Food	
Eating out & school lunch fees	Meals outside the home	
Rents, repairs & maintenance[108]		Rent + repairs & maintenance
Fuel, light & water charges	Fuel, light & water charges	
Furniture/household utensils /electric appliances[109]		Furniture & household utensils + durable goods assisting housework + heating & cooling appliances
Digital home appliances		Recreational durable goods - pianos - desks
Clothes & footwear	Clothes & footwear	
Medical care	Medical care	
Transportation		Public transportation + Private transportation
Communication	Communication	
Internet communication charge	Internet connection charge	
Education	Education	
Reading & recreation		Culture & recreation - recreational durable goods - internet connection charges
Other living expenditures	Miscellaneous	

[105] Questionnaire categories: Agriculture, forestry, fisheries; and Mining.

[106] Questionnaire categories: Sales worker e.g. storekeeper, real estate broker; Service worker e.g. barber, hotel employee; Transportation and communication worker e.g. cable operator, car/ship driver; and Preservation/guards worker e.g. firefighter, police officer.

[107] Questionnaire categories: Administrator e.g. congressman, manager; Office worker e.g. accountant; Information processing engineer e.g. computer programmer; and Professional or technological worker e.g. engineer, teacher.

[108] Excluding housing loan and common-area charge of apartment.

[109] Excluding digital appliances.

Table 2.23: Using subjective well-being measures as a proxy for utility: Examples from the literature

Authors	Area	Type of Data	Object of interest	Subjective well-being measure	Estimation of object of interest	Policy interpretation
Di Tella, MacCulloch, Oswald (2001)	Macroeconomics	Cross-country, individual-level panel data, OECD macroeconomic data	MRS between inflation and unemployment	Overall life satisfaction	Ratio of coefficients	Optimal inflation - unemployment tradeoff in monetary policy
Di Tella, MacCulloch, Oswald (2003)	Macroeconomics	Cross-country, individual-level panel data, OECD macroeconomic data	MRS between GDP per capita and unemployment	Same as Di Tella et al (2001)	Ratio of coefficients	"Psychic" cost of recessions (equivalent monetary compensation for GDP declines or job loss)
Ferrer-i-Carbonell, van Praag (2002)	Health economics	Individual-level panel data	MRS between health satisfaction and income	Domain-specific life satisfaction	Ratio of coefficients	Monetary value of being chronically ill (compensating and equivalent variation)
Finkelstein, Luttmer, Notowidigdo (2013)	Health economics	Individual-level panel data	Effect of chronic disease on the marginal utility of (non-medical) consumption	Overall life satisfaction	Reduced-form estimation of life cycle model with CRRA utility function; ratio of coefficients between health status and income	Optimal level of health insurance benefits and optimal level of life cycle savings
Frijters, Johnston, Shields (2011)	Health economics	Individual-level panel data	Income required to compensate for a major life event	Overall life satisfaction	Per-period compensation that equates the discounted life satisfaction of the individual if the event had not occurred; as well as standard valuation approach	Accounting for dynamic effects of life events when calculating appropriate compensation
Groot, Maassen van den Brink, Plug (2004)	Health economics	Individual-level panel data	Equivalent variation for cardiovascular disease	Income evaluation question	Per-period compensation that equates the lifetime utility of an individual if the event had not occurred	Cost-benefit analysis of health interventions
van den Berg, Ferrer-i-Carbonell (2007)	Health economics	Individual-level cross sectional data	MRS between income and hours of informal care	Overall happiness	Ratio of coefficients	Compensating variation for an hour of informal care
Levinson (2009)	Environmental economics	Individual-level cross-sectional surveys, EPA air quality data	Average MRS between income and pollution	Overall happiness	Ratio of coefficients	Willingness to pay for improved air quality
Luechinger (2009)	Environmental economics	Individual-level panel data; SO_2 monitor-level data	Tradeoff between pollution and income (compensating surplus)	Overall life satisfaction	Ratio of coefficients	Total willingness to pay for improvements in air quality, cost of migration
Plug, van Praag (1995)	Welfare economics	Individual-level cross sectional data	Elasticity of welfare with respect to family size (shadow price of a child)	Cantril question and income evaluation question	Ratio of coefficients	Constructing equivalence scales
Layard, Mayraz, Nickell (2008)	Public economics	Individual-level cross-sectional and panel surveys from different countries	Elasticity of marginal utility of income (with respect to income)	Overall happiness or life satisfaction	Maximum likelihood estimation of a happiness-income function, assuming CRRA utility	Cost-benefit analysis, optimal taxation
Oswald, Powdthavee (2008)	Tort law	Individual-level panel data	MRS between income and life event	Overall life satisfaction	Ratio of coefficients	Appropriate compensation for emotional damages

Appendix 2B: Additional GARP tests using imputed happiness measures

This section presents further attempts to address the issues with the happiness measures discussed in Section 2.3.1, which provide a robustness check for the unitary model results and demonstrate that reported happiness is generally inconsistent with the data.

To address the use of respondent's happiness only, a more aggregate happiness measure can be constructed using either the average or the total of the respondent and spouse's overall happiness (domain-specific happiness is unreported for most spouses). This approach requires interpersonal comparability of happiness reports, but is the most tractable way to aggregate happiness reports for use in the unitary model.

To address the domain mismatch, consumption-specific happiness can be imputed from the other domain-specific ratings by assuming that overall happiness is the average or weighted average of happiness over all domains. This approach is motivated by the two-layer model proposed by van Praag et al (2003), in which overall life satisfaction is an additively separable function of all domain-specific satisfactions[110]. Under the assumption that the overall happiness rating H_t is a simple average of D domain-specific happiness ratings $\{H_t^1, ..., H_t^D\}$, unobserved consumption-specific happiness (H_t^D) can be calculated from the $D-1$ observed domains:

$$H_t^D = DH_t - \sum_{d=1}^{D-1} H_t^d$$

Figure 2.11 shows the distribution of imputed consumption happiness. Around 30% of the imputed values were inside the survey range (0 to 10) and could therefore constitute a valid survey response. This method can give imputed consumption happiness ratings that are

[110] There is some evidence from the JHPS data used in this chapter to support a linear relationship between overall and domain-specific happiness. In a pooled linear regression of overall happiness on domain-specific measures, all coefficients except the one on leisure time (amount) are significant.

outside the survey range. For example, if overall happiness was rated 5 but domain-specific happiness was 10, then the unobserved consumption happiness must be negative (-30 if there are 8 domains total) in order to bring the simple average down to 5.

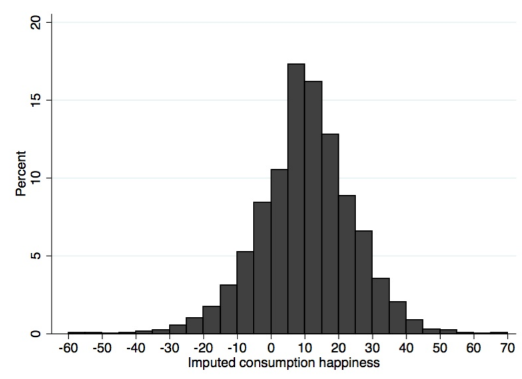

Figure 2.11: Imputed consumption happiness rating (simple average)

The assumption that individuals weight all domains equally may not be plausible. Individuals may place greater weight on certain domains according to their preferences, for example a workaholic might put higher weight on happiness from employment and lower weight on happiness from leisure time. Unlike a simple average, this assumption does not lead to an obvious unique solution, because in each period there are $D+1$ unknowns (the weakly positive weights on each domain ($w_t^1 \geq 0, \ldots, w_t^D \geq 0$) and the consumption happiness H_t^D), but only 2 equations that relate them[111]. Even the extra condition that H_t^D must be an integer between 0 to 10 does not help narrow down the possible combinations.

Since it is unlikely that any of the domains has no effect on overall happiness, a plausible requirement for choosing these weights is to 'spread them out' over all the domains as much as possible, so that each domain makes some contribution to overall happiness. One way to formalise this idea is to use concepts from information theory. Thinking of these weights as representing probabilities of outcomes and each domain as a possible outcome, 'spreading out' corresponds to choosing the weights that maximise the uncertainty of the final out-

[111] 1. The weighted sum must equal overall happiness ($\sum_{d=1}^{D} w_t^d H_t^d = H_t^D$), and 2. the weights must sum to 1 ($\sum_{d=1}^{D} w_t^d = 1$).

come (otherwise known as **entropy**). Figure 2.12 gives an example of a high-entropy and low-entropy set of weights. The high-entropy set allows each domain to contribute towards the average whereas the low-entropy set puts all the weight on one domain. Using the probability analogy, in the left panel each domain has a similar probability of occurring so the outcome from a random draw would be very 'uncertain' (entropy is high), but in the right panel entropy is low because the outcome is certain.

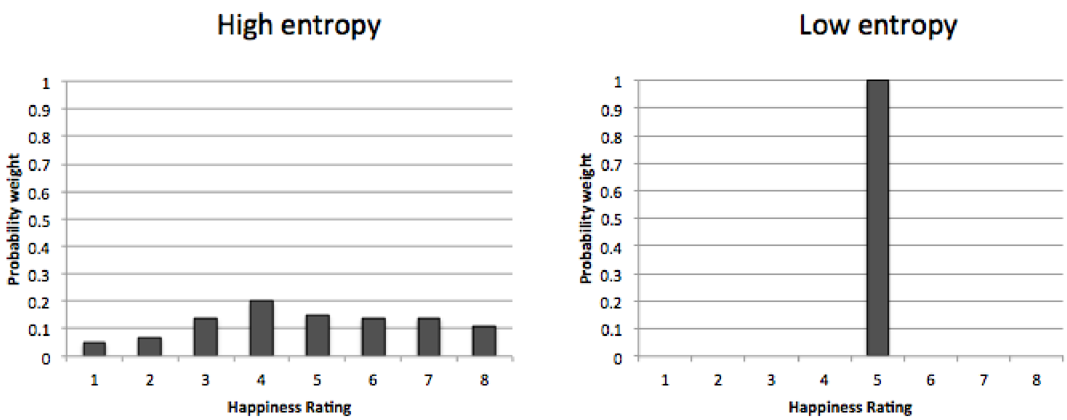

Figure 2.12: High vs. low entropy

Entropy is measured as a function of the weights and is defined as:

$$H(\mathbf{w}) = -\sum_{i=1}^{D} w_t^i log(w_t^i)$$

If all happiness ratings H_t and $\{H_t^1, ..., H_t^D\}$ were known, then without additional information or assumptions about the weights, the maximum entropy estimates are the weights $\mathbf{w} = \{w_t^1, ..., w_t^D\}$ on the domain-specific happiness ratings that maximise this entropy function. For example, without the range restriction on H_t^D, entropy is maximised when all probabilities are equal (a simple average), since there is no reason to place more weight on one particular domain. However, consumption-specific happiness H_t^D is unknown so must be found along with the set of weights.

Since consumption-specific happiness can only take a limited number of values, the set of weights and consumption-specific happiness that jointly maximise entropy can be found using the following 2-step process: For every integer from 0-10, 1. Find the weights that maximise $H(\mathbf{w})$, 2. Evaluate the entropy function $H(\mathbf{w})$ at the optimal weights. The imputed consumption happiness value H_t^D is the integer that gives the highest value of $H(\mathbf{w})$.

Figure 2.13 shows the imputed values of consumption happiness selected using maximum entropy, and the corresponding probability weight on the unobserved consumption domain. As mentioned earlier, maximising entropy involves choosing a simple average whenever possible. Since most imputed values of H_t^D were outside the survey range, most entropy-maximising values were at the extreme ends of the scale, resulting in a vastly different distribution compared to those in Figures 2.5 and 2.6. Although a positive weight was always assigned to consumption happiness, consumption did not count towards overall happiness for a large proportion of individuals because their assigned happiness rating was 0. The spike at 0.12-0.13 (around $\frac{1}{8}$) represents the values of H_t^D that were already inside the survey range when a simple average was taken.

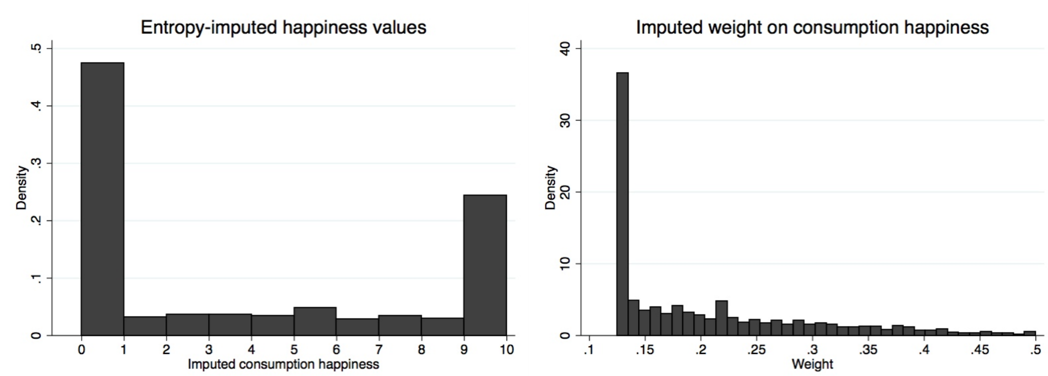

Figure 2.13: Imputed consumption happiness values and weights

The GARP test was conducted using the average or total of respondent's and their spouse's happiness (GARP-H (avg) and GARP-H (total), respectively). The GARP test was also repeated for consumption-specific happiness (GARP-HD), which was imputed from the respondent's ratings over other domains[112]. Compared to GARP-H, pass rates were almost

[112] As mentioned in Section 2.3.1, Spouse's ratings for domain-specific happiness were often missing (not available for more than one year, or not all domains were rated), so could not be used.

twice as high (15-20% compared to 8-10%), but are still much lower than the standard model. Also, as discussed in Section 2.3.1 over half of the imputed values were outside the scale given in the survey (0-10) so cannot be interpreted as a valid survey response.

The GARP test was also repeated using the values of H_t^D selected via maximum entropy (GARP-E), as outlined in Section 2.3.1. Unlike the simple average method, these imputed ratings of consumption happiness were constrained to be within the happiness scale given in the survey. However, this range restriction resulted in pass rates similar to that of GARP-H, which suggests that the much lower pass rates compared to GARP are not due to a domain mismatch. Table 2.24 lists all the revealed preference tests conducted in this chapter, and Table 2.25 summarises the results of all GARP tests for the unitary model. These results suggest that happiness measures are generally inconsistent with expenditure data, regardless of the measure used.

Table 2.24: Summary of revealed preference tests: Unitary model

Model	Happiness measure used	Abbreviation
Unitary	None	GARP
	Respondent's happiness	GARP-H
	Respondent and spouse's happiness (average or total)	GARP-H (avg), GARP-H (total)
	Imputed consumption happiness (simple average)	GARP-HD
	Imputed consumption happiness (entropy)	GARP-E

Table 2.25: Pass rate comparison (2011-2013, standard errors in brackets): Unitary test results

	All households	1-person	2-person
GARP	0.987	0.994	0.987
	(0.113)	(0.077)	(0.115)
GARP-H (overall)	0.086	0.101	0.091
	(0.280)	(0.303)	(0.288)
GARP-H (income)	0.090	0.114	0.096
	(0.286)	(0.318)	(0.295)
GARP-H (avg)	0.107	0.107	0.120
	(0.310)	(0.309)	(0.325)
GARP-H (total)	0.108	0.107	0.122
	(0.311)	(0.309)	(0.327)
GARP-HD	0.154	0.201	0.164
	(0.361)	(0.402)	(0.370)
GARP-E	0.072	0.084	0.064
	(0.258)	(0.278)	(0.245)

3. Price inattention, inflation misperceptions, and cost-of-living indices

3.1 Introduction

There is survey evidence that public perceptions of inflation differ from official statistics such as the CPI and RPI (O'Donoghue, 2007; Bank of England, 2010)[113]. For example, according to the Bank of England survey, perceived inflation in the UK from 2000-2010 has generally been higher than actual CPI inflation and loosely tracks RPIX inflation (Figure 3.14 below).

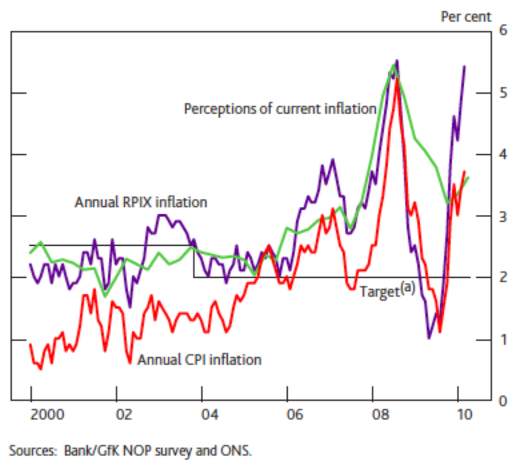

Figure 3.14: Actual and perceived inflation (Bank of England, 2010)

There are several reasons why perceived inflation may differ from official inflation statistics.

[113]Examples of these surveys include the Bank of England/Gfk NOP Inflation Attitudes survey, which has been conducted regularly since 1999, and the Joint Harmonised EU Programme of Business and Consumer Surveys. Respondents are asked questions such as "How do you think that consumer prices have developed over the last 12 months?" or "Do you believe prices have increased or decreased in the past 12 months?", and choose their answer from the following options (wording may differ slightly across surveys): risen a lot, risen moderately, risen slightly, stayed about the same, fallen, don't know. Perceived inflation is taken as the percentage of survey respondents who believed that prices have increased.

Measures such as the CPI and RPI are based on behaviour of the 'average household', which could differ greatly from expenditure patterns of individual hosueholds. Respondents may not fully understand how official statistics are calculated or how to relate these statistics to their personal experiences[114]. Individuals may also perceive inflation differently due to behavioural biases, demographic characteristics, or purchasing behaviour. Studies on inflation perceptions have found that individuals tend to notice price increases more than price reductions (Brachinger, 2006), and notice price changes for items they buy more frequently (Antonides et al, 2006). High-income earners also tend to have lower perceptions of inflation rates compared to low-income earners (Biau et al, 2010). Attention to prices is an important factor in correctly perceiving inflation: as shown in Figure 3.15, perceived inflation is correlated with food and energy prices, which consumers encounter more frequently than other prices and are therefore more salient, as well as news coverage of inflation. While the effect of inflation misperceptions on monetary policy and inflation expectations have been studied extensively, the implications of price inattention on perceived inflation and welfare analysis have not yet been explored in depth.

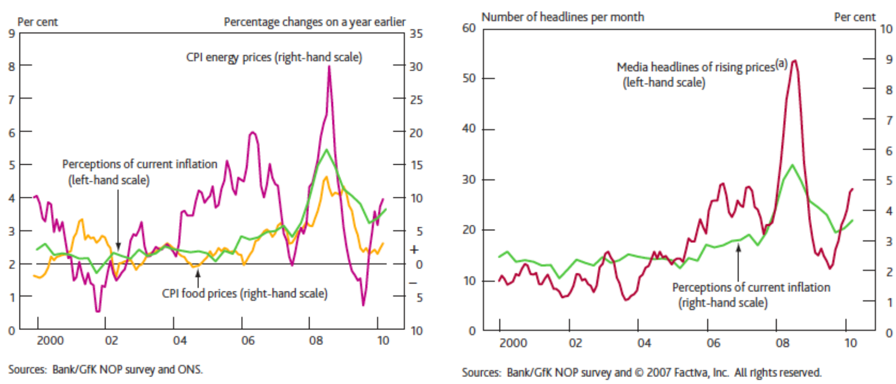

Figure 3.15: Perceived inflation, food and energy prices, and news coverage (Bank of England, 2010)

[114]To address this issue, the ONS launched the Personal Inflation Calculator in 2007, aimed at improving public confidence in official statistics and knowledge of how inflation is calculated. Also, the European Commission lists the inflation rates for frequently purchased goods (bought on a monthly basis or more often) alongside the official inflation rate, so consumers can better understand why overall inflation may not coincide with price changes in the goods they typically encounter.

This chapter explores the effect of price inattention on the perceptions of inflation and the cost of living, using a behavioural model in which consumers notice 'large' price changes but may not notice 'small' price changes. 'Large' and 'small' are defined according to a price cutoff, which represents the smallest price change that a consumer notices and hence the degree of attention paid to prices. This cutoff can be specific to an individual and a particular good. This model of price inattention is characterised by short-term memory, an asymmetric response to overspending and underspending, and price cutoffs. These features capture the notions of anchoring, partial adjustment towards the truth, and framing effects, which previous studies have shown to be part of behavioural decision-making[115].

This chapter then derives necessary and sufficient conditions for data to be consistent with this price inattention (PI) model, which are easy to test using observable information only. The revealed preference test is a simple extension of the standard test for static utility maximisation given in Varian (1982), with perceived prices as an additional unknown that must be found.

A data application, using the Stanford Basket Dataset, demonstrates that this model has some desirable properties: (1) it is plausible, meaning that the a considerable proportion of household behaviour can be rationalised by assuming a common cutoff for all goods, and the minimum cutoff found for households falls within a reasonable range; and (2) it is falsifiable, so even though perceived prices are unknown, it is not always possible to find a set of perceived prices that can rationalise the given behaviour. By comparing perceived cost-of-living indices and annual inflation rates with their standard counterparts, this model also generates the intuitive result that consumers have more accurate perceptions of inflation and generally behave neoclassically when inflation is high, but may substantially misperceive inflation when it is low. This result has important implications for welfare analysis, as well as

[115]See Tversky and Kahnemann (1975) and Tversky and Kahnemann (1985) for further details and specific examples of behavioural decision-making.

measurements of the cost of living, when consumers are not fully attentive to price changes.

3.1.1 Modelling price inattention

Inattention has been modelled in stochastic choice settings: for example, in the Fuzzy Attention Model of Aguiar (2015), decision makers have rational preferences but do not pay full attention to all items in their choice set, leading to violations of the Independence of Irrelevant Alternatives (IIA). The decision maker has an 'attention capacity' for each item that depends on the complementarity or substitutability of the other items in the set, and items added to the set. Aguiar (2015) develops revealed preference conditions for consistency with this model, which use the violations of IIA to infer the underlying preference relation, though these have not yet been tested on data.

The price inattention literature most closely related with the themes of this chapter is the 'sparse-max' model of Gabaix (2014), in which the consumer maximises utility under perceived prices that are formed using a subset of available information. The degree of attention paid to each piece of information is chosen optimally according to the cost of solving a more complicated problem and the benefits of making better choices with that information. These are some theoretically appealing aspects of the model, but make it untestable using observable data only. However, this chapter presents a version of this model that is testable using data on prices and quantities only.

This chapter relaxes one assumption about neoclassical consumer behaviour, which is that consumers notice and optimise according to the actual price of goods. The behavioural consumer in this chapter does not notice the actual price of all goods unless the price difference from last period is 'large enough', and instead maximises utility according to his perceived prices, which may not always be the same as the actual prices if period-to-period price changes are too 'small' to notice.

Behaviour according to vague descriptors such as 'small' and 'large enough' can be modelled using concepts from fuzzy set theory, which deals with non-sharp boundaries of classification, in which membership of a particular set is a matter of degree rather than a binary characteristic[116]. Fuzzy sets are characterised by a continuous membership function with bounded support, which determines the degree ($\alpha \in [0, 1]$) to which a given element in the domain belongs to that set[117].

In the context of price inattention, the price of good k in period $t - 1$ (p_{t-1}^k) is part of a fuzzy set, which contains all prices considered to be 'essentially the same as p_{t-1}^k'. Each fuzzy set is characterised by a convex and continuous membership function ($\alpha(\triangle p_k)$) of bounded support that determines the degree ($\alpha \in [0, 1]$) to which a given price belongs to that set, with $\alpha = 1$ at p_{t-1}^k. Only prices with a 'degree of belonging' of 0 are definitely not part of the set; all prices with $\alpha > 0$ are part of the set to some degree. As shown in Figure 3.16, there are a wide variety of membership functions that satisfy these properties, including standard kernel functions[118].

[116] Zadeh (1965) introduced fuzzy sets to mathematically represent information or data that had nonstatistical imprecision or uncertainty. Fuzzy logic, which is derived from fuzzy set theory, is used in computing contexts to represent the approximations used in human reasoning and for generating discrete outputs from a continuous input (for example, classifying someone as 'young' based on their age, or classifying a good as 'cheap' based on its price).

[117] Formally, if B is a nonempty set, then a fuzzy set $B^{'} \subseteq B$ is characterised by its membership function $\alpha_{B'} : B \rightarrow [0, 1]$, where $\alpha_{B'}(b)$ denotes the degree of membership of each element $b \in B$.

[118] Although the range of α suggests that it can be interpreted as a probability (the likelihood that two numbers are similar) or a weight (as in kernel density estimation), the analogy is not exact as membership functions do not have to have the same properties as probability density functions or kernel functions (area underneath does not have to sum to 1), unless membership functions are rescaled accordingly.

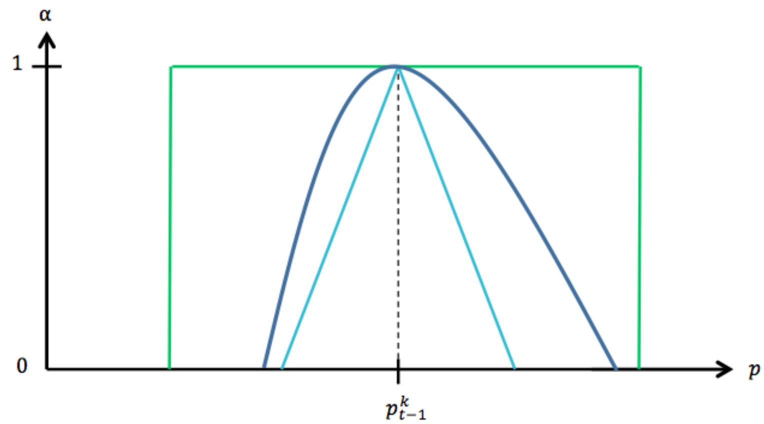

Figure 3.16: Examples of membership functions

This formulation has similar themes to the behavioural economics research on limited attention, in which decision makers may choose not to consider all available options or all aspects of each option. Models of limited attention usually contain an attention parameter that ranges from 0 to 1, which determines the extent to which the decision-maker perceives or notices certain aspects of his choice problem. For example, in the discrete choice setting of Aguiar (2015), the decision maker assigns an 'attention capacity' to each alternative in his choice set, and in the standard consumer theory setting of Gabaix (2014), the consumer chooses how much attention to pay to each price. In the price-inattention model described in Section 3.1.1 and , the α parameter can be interpreted as the degree of attention paid to the current price of a good, which in this case is either complete attention or no attention.

This chapter focuses on price perceptions characterised by the symmetric rectangular membership function, shown by the uniform function in Figure 3.16. Under this function, the consumer only notices a price change if it is 'large enough' in absolute terms, otherwise he doesn't notice the change and optimises according to the previous price. To give a concrete example, suppose that $p^k_{t-1} = 1.00$, and the consumer only notices price changes that are greater than or equal to 5 cents. In this case, α is an indicator function that equals 1 if $0.95 < p^k_t < 1.05$ (consumer does not notice the price difference) and $\alpha = 0$ otherwise (consumer notices the price difference). In general, the consumer's perceived price for good k in period t ($\tilde{p}_t^k$) can be written as

$$\tilde{p}_t^k = \alpha_t^k p_{t-1}^k + (1 - \alpha_t^k) p_t^k \tag{1}$$

where $\alpha_t^k = 1$ if $|p_t^k - p_{t-1}^k| \leq c^k$, a good-specific cutoff, and $\alpha_t^k = 0$ otherwise. This expression is similar to the one derived in the 'sparse-max' model of Gabaix (2014), but as Section 3.1.2 explains, the other aspects of the decision-making process lead to different observable behaviour. Throughout this chapter, consumers who perceive prices in this way will be referred to as 'price-inattentive' consumers, and their behaviour may be referred to as 'fuzzy'[119].

The size of the cutoff determines the degree of price inattention. For computational purposes, the cutoff is assumed to be time-invariant for a particular good bought by a given consumer[120]. This chapter will also remain agnostic about how the cutoff is determined (see Section 3.1.2 for further discussion), though intuitively, the cutoff would likely depend on the price range of the good or characteristics of the consumer. For example, if a product became $1 cheaper today, the consumer might consider the new price 'essentially the same' as the old price if the product was a bed, but not if it was a pencil. Richer consumers may also have larger cutoffs because they can afford to pay less attention to prices.

In the one-good case, the price-inattentive consumer only behaves like a neoclassical consumer if the price difference passes a certain threshold ($\alpha = 0$ for the new price). Otherwise, he optimises according to last period's price. In the multi-good case, the price-inattentive consumer's behaviour only coincides with the neoclassical consumer's if he correctly perceives the price changes of all goods in each period. Otherwise, his perceived prices, and hence the price ratio between any two goods, may not equal the observed prices or price ra-

[119] Note that this terminology is based on the notion of fuzzy sets, and does not refer to statistical concepts such as measurement error.

[120] The existence of necessary and sufficient conditions for rationalisability do not depend on this assumption, but less restrictive assumptions on the cutoff would be computationally burdensome. Section 3.3.1discusses the computational issues in more detail.

tios. His deviation from neoclassical behaviour therefore arises due to misperceived relative prices. Conditioning behaviour on last period's prices introduces another departure from neoclassical behaviour: path dependence. It is possible that a consumer may choose differently when the same price change is presented as a large change compared to a series of very small changes (similar to a framing effect).

This model also incorporates other aspects of behavioural decision-making, such as anchoring, partial adjustment towards the truth, and framing effects (Tversky and Kahnemann, 1975). The consumer's perceptions of this period's price are anchored on last period's price, and will not be correct unless all price changes are large enough or all prices have not changed since last period. Otherwise, perceived prices will be a combination of current and last period's prices. The consumer's final response to a series of small price changes (which he does not notice) can also be different to his response to one large price change that equals the sum of these small changes (which he does notice).

3.1.2 Links to the Gabaix (2014) 'sparse-max' model

In Gabaix (2014), the consumer chooses according to perceived prices, but faces a hard budget constraint determined by actual prices and his perfectly-perceived per-period income x_t, so he may face discrepancies between his expected and actual expenditure. For example, if the prices of all goods decrease but are not perceived, the consumer underspends. Alternatively, if the prices of all goods increase but are not perceived, the consumer overspends. In the simplest model of behaviour, per-period utility maximisation with no borrowing or saving, overspending is not possible so the consumer must choose another bundle. In Gabaix (2014), the consumer reacts symmetrically to overspends and underspends by adjusting his perceived expenditure so that he exhausts his budget constraint under the misperceived prices[121].

Since this adjustment is unobservable by the researcher, expenditure is taken to be exogenous in this chapter, and the consumer solves a 'perceived' cost-minimisation problem[122]. By duality with utility maximisation, this problem is equivalent to that presented in Section 3.2.

This failure to completely adjust to all price changes has observable implications that only require data on prices and quantities, as in the test for static utility maximisation (GARP test) presented in Varian (1982)[123]. Figure 3.17 shows how rational choices (from the consumer's viewpoint) made based on perceived instead of actual prices can violate GARP. In period 2, both price changes were not noticed so the consumer perceived a parallel shift of his period 1 budget constraint and made a choice that appeared rational, but this choice violated GARP according to the actual budget constraint (observed by the researcher). The different slopes of the actual and perceived budget constraints in period 2 illustrates the fact that violations are due to misperceived relative prices.

[121] Formally, the consumer maximises according to a 'perceived budget constraint' (expenditure level $\tilde{x}_t$), where $\tilde{x}_t$ is chosen such that the observed demands under misperceived prices ($\mathbf{q}_t(\tilde{\mathbf{p}}_t, \tilde{x}_t)$) are optimal, and the actual budget constraint is exhausted ($\mathbf{p}_t \cdot \mathbf{q}_t(\tilde{\mathbf{p}}_t, \tilde{x}_t) = x_t$). Graphically, this adjustment is a parallel shift of the consumer's misperceived budget constraint until it intersects the actual budget constraint at the point where his indifference curve is tangent to the misperceived budget constraint (see Figure II in Gabaix (2014)).

[122] That is, min $\tilde{\mathbf{p}} \cdot \mathbf{q}_t$ s.t. $u(\mathbf{q}_t) \geq u_t$, where u_t is a minimum level of utility that the consumer desires to reach.

[123] This test checks whether or not there exists a rational (complete and transitive) preference relation over bundles of goods that can explain the observed behaviour. If such a preference relation exists, then the data is said to satisfy GARP, and is consistent with utility maximisation. According to Afriat's Theorem, data is also consistent with static utility maximisation if and only if it satisfies the following set of linear inequalities: $u_s \leq u_t + \lambda_t \mathbf{p}_t \cdot (\mathbf{q}_s - \mathbf{q}_t)$ for all $s, t \geq 2$, which will henceforth be referred to as 'Afriat inequalities' (Varian, 1982).

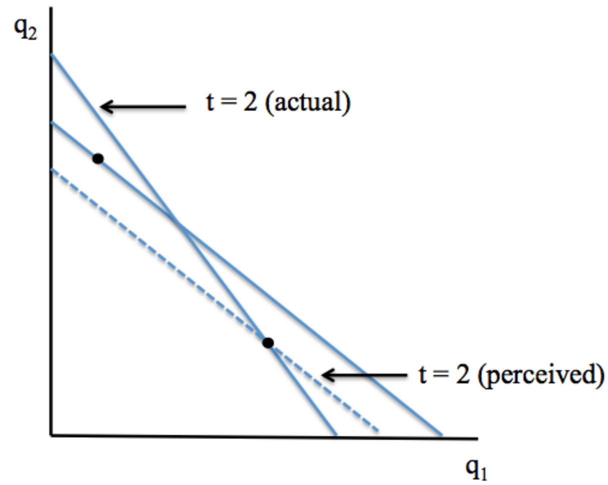

Figure 3.17: "Rational" GARP violation

As mentiond in Sections 3.1.1 and 3.1.2, this behavioural model is a special case of the 'sparse-max' model of Gabaix (2014), with the following simplifications:

1. Default price: The default price in Gabaix (2014) is the long-run average price, which implies that the consumer has a long-term memory. However, it is difficult to construct this default price, since each consumer may have a different notion of what 'long-term' is, or may use some other heuristic to determine what the default is (for example, rounding to the nearest whole number). In my model, consumers have short-term memory (they only remember last period's prices) so the default price is last period's price, which makes the model tractable and testable. This short-term memory results in path dependence, which is an interesting and realistic element of my model.

2. Determinants of the attention parameter: Gabaix (2014) explicitly models the choice of the attention parameter, which is determined by the optimal tradeoff between adding more elements to the decision-making process and the utility losses from sticking to the default option. My model remains agnostic about how the attention parameter (the cutoff) is determined, so while it is possible that the consumer does optimise ac-

cording to a two-stage decision process (choose the optimal cutoff, then choose what to purchase), the first stage is unobservable and the revealed preference test can only be based on what we observe from the second stage (prices and quantities). Section 3.3.3 explores the relationship between the cutoff size and variables that could affect the degree of price misperceptions, such as income and family size.

3.2 Revealed preference conditions

In each period, the price-inattentive consumer solves the standard static maximisation problem, but with an additional intertemporal condition on prices in adjacent periods:

$$\max u(\mathbf{q}_t) \text{ s.t. } \tilde{\mathbf{p}}_t \cdot \mathbf{q}_t = x_t$$

$$\tilde{p}_t^k = \alpha_t^k p_{t-1}^k + (1-\alpha_t^k)p_t^k \text{ for } t = 2, ..., T \text{ and } k = 1, ..., K$$

$$\alpha_t^k = 1 \text{ if } |p_t^k - p_{t-1}^k| \leq c^k, \text{ and } \alpha_t^k = 0 \text{ otherwise, for } t = 2, ..., T \text{ and } k = 1, ..., K \tag{R}$$

where $u(\mathbf{q}_t)$ is a well-behaved utility function. Under this formulation, the consumer only optimises according to the new price if the price difference is 'large enough' i.e. larger than a good-specific cutoff value c^k. This cutoff only depends on the absolute value of the price difference, meaning that the consumer evaluates price increases and decreases in the same way. The cutoff is defined in terms of absolute price differences in order to be easily interpretable, but it is also possible to define the cutoff as a perentage of the price change without fundamentally changing the revealed preference conditions and perceived prices, which only depend indirectly on the cutoff through α_t^k[124].

[124]Condition (R) then becomes: $\alpha_t^k = 1$ if $\frac{|p_t^k - p_{t-1}^k|}{p_{t-1}^k} \leq c^k$, and $\alpha_t^k = 0$ otherwise, for $t = 2, ..., T$ and $k = 1, ..., K$.

In terms of modelling how prices are perceived, defining cutoffs in levels is more plausible than percentages because consumers are more likely to remember the price level of the products they purchase and hence evaluate changes in terms of levels rather than percentage changes. Using levels is more plausible when product prices are similar, as in the supermarket shopping context of Section 3.3. When product prices are in different orders of magnitude (for example, household expenditure data on electronics, stationary, and clothing), then percentages may be more plausible.

The Lagrangian is $\mathcal{L} = u(\mathbf{q}_t) + \lambda_t(x_t - [\boldsymbol{\alpha}_t\mathbf{p}_{t-1} + (1-\boldsymbol{\alpha}_t)\mathbf{p}_t]\cdot\mathbf{q}_t)$, and the first-order conditions are:

$$\frac{\partial u(\mathbf{q})}{\partial q_t^k} = \lambda_t[\alpha_t^k p_{t-1}^k + (1-\alpha_t^k)p_t^k] = \lambda_t \tilde{p}_t \text{ for } t = 2, ..., T \text{ and } k = 1, ..., K \qquad \text{(FOC)}$$

Therefore by concavity of the utility function, the Afriat inequalities are:

$$u_s \leq u_t + \lambda_t \tilde{\mathbf{p}}_t \cdot (\mathbf{q}_s - \mathbf{q}_t) \text{ for all } s, t \geq 2$$

These inequalities are non-linear in unknowns due to interactions between λ_t and α_t^k. However, these inequalities are linear conditional on the cutoff vector $\mathbf{c} = [c_1, ..., c_k]$, which in this case directly determines both $\boldsymbol{\alpha}_t$ and $\tilde{\mathbf{p}}_t$. For the rectangular membership function, the cutoff indicates which price changes are perceived and hence whether the consumer acts according to last period's price (p_{t-1}^k and $\alpha = 1$) or this period's price (p_t^k and $\alpha = 0$). Finding cutoffs that generate $\boldsymbol{\alpha}_t$ and $\tilde{\mathbf{p}}_t$ that are consistent with the data is less computationally intensive than directly finding a combination of $\{\alpha_t^k, p_{t-1}^k, p_t^k\}_{t=2,...,T}^{k=1,...,K}$ that is consistent with observed behaviour. The good-specific cutoffs $\mathbf{c}$ are unknowns, but can be found via grid searching over a plausible range[125].

Rationalisability for the price-inattention model can therefore be defined as follows:

Definition 3.1. The dataset $\mathbf{D} = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ is rationalisable by the price-inattention model (**pi-rationalisable**) if there exists a well-behaved utility function $u(\mathbf{q})$ and numbers

Another way to define perceived prices is in terms of the largest unit, which is known in the behavioural economics literature as the 'left-digit effect' (Basu, 1997; Stiving, 2000). For example, 9.99 is perceived as 9, whereas 10.01 is perceived as 10. In this case, perceived prices would be defined as $\tilde{p}_t^k = \lceil p_t^k \rceil$. However, it is more difficult to test for this type of perceptions using aggregated prices, or if most prices ended in 9 because sellers recognised this type of bias in consumers (in the dataset used for this chapter, 58% of all disaggregated prices ended in 9).

Finally, note that the cutoff does not have to be the same for price increases and decreases, but we impose symmetry to lessen the computational burden.

[125] To lessen the computational burden, the restriction that $c^k = c$ for all k can also be added, which is plausible if the price ranges of the goods are similar. For example, a consumer may have a similar cutoff for pens and pencils, but not chairs and beds. The fineness of the grid is naturally limited by the smallest unit of currency e.g. 1 cent, since a consumer who has a cutoff smaller than this would be observationally equivalent to a neoclassical consumer.

Also, the upper bound for the common cutoff could be set at the largest price difference between adjacent time periods i.e. $c = max(|\Delta p_t^k|)$ for $\forall k$; $t \geq 2$, since any larger cutoff would generate the same set of $\boldsymbol{\alpha}_t$ and $\tilde{\mathbf{p}}_t$.

$\lambda_t > 0$ and $\alpha_t^k \in \{0, 1\}$ that satisfy the first order conditions (FOC) and the restriction (R) on the membership function $\alpha^k(\triangle p)$ for all t, k.

The following theorem uses the link between the good-specific cutoffs and the indicator function $\alpha^k(\triangle p)$ to derive necessary and sufficient conditions for this price-inattention model.

Theorem 3.1. The following statements are equivalent:

(1) The dataset $\mathbf{D} = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ is pi-rationalisable by a symmetric, good-specific rectangular membership function $\alpha^k(\triangle p)$.

(2) There exist numbers $u_t \in \mathbb{R}$, $\lambda_t > 0$, and $c^k \geq 0$ such that for $t = 2, ..., T$ and $k = 1, ..., K$

$$u_s \leq u_t + \lambda_t \tilde{\mathbf{p}}_t \cdot (\mathbf{q}_s - \mathbf{q}_t)$$

$$\tilde{p}_t^k = \alpha_t^k p_{t-1}^k + (1 - \alpha_t^k) p_t^k$$

$$\alpha_t^k = 1 \text{ for } \left|\triangle p_t^k\right| \leq c^k \text{ and } \alpha_t^k = 0 \text{ otherwise.}$$

***Proof*:**

(1) $\implies$ (2) [*Necessity*] The dataset $\mathbf{D} = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ is pi-rationalisable by a symmetric, good-specific rectangular membership function if there exists a locally non-satiated, continuous, concave, and differentiable utility function $u(.)$ and numbers $\lambda_t > 0$, $c^k \geq 0$, and $\alpha_t^k \in \{0, 1\}$ such that for $t = 2, ..., T$ and $k = 1, ..., K$

$$\frac{\partial u(\mathbf{q})}{\partial q_t^k} = \lambda_t[\alpha_t^k p_{t-1}^k + (1 - \alpha_t^k) p_t^k]$$

$$\tilde{p}_t^k = \alpha_t^k p_{t-1}^k + (1 - \alpha_t^k) p_t^k$$

$$\alpha_t^k = 1 \text{ for } \left|\triangle p_t^k\right| \leq c^k \text{ and } \alpha_t^k = 0 \text{ otherwise}$$

By concavity of $u(.)$,

$$u_s \leq u_t + \nabla u(\mathbf{q}_t)'(\mathbf{q}_s - \mathbf{q}_t) \text{ for } \forall s, t = 2, ..., T$$

Substituting $\lambda_t[\boldsymbol{\alpha}_t \mathbf{p}_{t-1} + (1 - \boldsymbol{\alpha}_t)\mathbf{p}_t]$ for $\nabla u(\mathbf{q}_t)$ gives statement (2):

$$u_s \leq u_t + \lambda_t \tilde{\mathbf{p}}_t \cdot (\mathbf{q}_s - \mathbf{q}_t)$$

$$\tilde{p}_t^k = \alpha_t^k p_{t-1}^k + (1-\alpha_t^k) p_t^k$$

$$\alpha_t^k = 1 \text{ for } \left|\triangle p_t^k\right| \leq c^k \text{ and } \alpha_t^k = 0 \text{ otherwise}$$

for $t = 2, ..., T$ and $k = 1, ..., K$.

(2) $\implies$ (1) [*Sufficiency*] By the definition of pi-rationalisable, we need to be able to construct the first-order conditions and the good-specific membership functions using the solution to the system given by (2).

Suppose there exist numbers $u_t \in \mathbb{R}$, $\lambda_t > 0$, and $c^k \geq 0$ such that

$$u_s \leq u_t + \lambda_t \tilde{\mathbf{p}}_t \cdot (\mathbf{q}_s - \mathbf{q}_t) \text{ for all } s, t \geq 2$$

$$\tilde{p}_t^k = \alpha_t^k p_{t-1}^k + (1-\alpha_t^k) p_t^k \text{ for all } t, k \geq 2$$

$$\alpha_t^k = 1 \text{ for } \left|\triangle p_t^k\right| \leq c^k \text{ and } \alpha_t^k = 0 \text{ otherwise}$$

Note that the cutoff determines α_t^k, which can then be used to construct the perceived prices $\tilde{p}_t^k$ from the observed prices in adjacent periods (p_{t-1}^k and p_t^k) for all $t, k \geq 2$. Given $\tilde{p}_t^k$, the first condition becomes the standard Afriat inequalities. By Afriat's Theorem, the dataset $\mathbf{D'} = \{\tilde{\mathbf{p}}_t, \mathbf{q}_t\}$ satisfies GARP and therefore there exists a locally nonsatiated, concave, continuous utility function $u(.)$ that is monotonic in $\mathbf{q}_t$ such that $\nabla u(\mathbf{q}_t) = \lambda_t \tilde{\mathbf{p}}_t$. Substituting for $\tilde{\mathbf{p}}_t$ obtains the first-order conditions of the price-inattention model. $\square$

For data to be consistent with the price-inattention model, it is therefore sufficient to find support prices $\tilde{p}_t^k$ such that the data satisfies the inequalities given in (2) (according to the equivalence between these inequalities and GARP, (2) can therefore be referred to as 'PI-GARP'). The rectangular membership function restricts the support prices to being either last period's price or this period's price. Finding the combination of past and current prices that are consistent with the data appears computationally difficult, but becomes simpler by conditioning on the cutoff. Each cutoff generates only one possible combination of $\{\alpha_t^k, p_{t-1}^k, p_t^k\}_{t=2,...,T}^{k=1,...,K}$ and hence one set of support prices $\{\tilde{\mathbf{p}}_t\}_{t=2,...,T}$ that can be substituted into the inequalities in (2). Given the cutoff, the inequalities can be solved in the same way as the standard Afriat inequalities (finding $u_t \in \mathbb{R}$ and $\lambda_t > 0$ for $t \geq 2$). However, there can be many possible cutoffs that are consistent with the data, so the solution may not be unique.

3.3 Empirical application

The following application, using the Stanford Basket Dataset, demonstrates that this price inattention model has some desirable properties: (1) it is plausible, meaning that the minimum cutoff found for households falls within a reasonable range; and (2) it is falsifiable, so not all behaviour can be rationalised with a large enough cutoff. The price-inattention model also generates the intuitive result that consumers have more accurate perceptions of inflation when inflation is high, but may not perceive inflation accurately when it is low.

3.3.1 Data description

Price-inattentive consumers can still pass GARP, either because all price changes are large enough or because the departure from neoclassical behaviour is not large enough to violate GARP. Therefore, the only way to distinguish a price-inattentive consumer from a neoclassical one is to find behaviour that fails GARP but passes PI-GARP. The Stanford Basket Dataset was chosen because a large proportion of households were shown to violate GARP in previous studies[126]. This panel dataset contains the daily barcode-level expenditures of 494 urban households for a range of supermarket goods, from June 1991 (Month 1) to May 1993 (Month 24), as well as demographic information and product details. These transactions took place across 4 stores in the South West of Chicago, and were collected by Information Resources, Inc[127].

[126]For example, Echenique, Lee, and Shum (2012) find that 80% of households violate GARP.
[127]The raw data was downloaded from the following link:
http://people.hss.caltech.edu/~mshum/res.html.

The full dataset contains 208,780 observations, but following Echenique, Lee, and Shum (2012), a subset of 13 food categories were used: bacon, barbecue sauce, cereal, chips/crisps, coffee, crackers, eggs, ice cream, nuts, pills, pizza, and sugar (108,201 observations total). Each observation records the household ID, item UPC code, day and week of purchase, number of units purchased, shelf price per unit, and the coupon value (if a household used a discount coupon). Shelf prices are used throughout this analysis, since incorporating coupon discounts into revealed preference analysis is problematic[128]. However, this is not a major limitation as most transactions (85%) occur at shelf prices.

Another limitation is that these prices are before-tax prices (tax is added at the checkout), and hence consumers may be violating GARP because they do not optimally account for the effects of taxes on their purchasing decisions[129]. However, it is difficult to disentagle the effects of tax salience and price attention without additional information or assumptions, and so we cannot address this issue with the available data.

The price and quantity data was aggregated by calendar month (24 months total) for each household, with prices weighted according to the number of units purchased[130]. Under this approach, households face different prices depending on their chosen product mix within each category. Using unit values as a price measure has some drawbacks, for example it cannot capture the effects of quality substitution, leading to biased estimates of price effects when used in demand systems[131] (Deaton, 1988).

[128]Specifically, we do not observe the potential discounts for goods that are not purchased, and cannot observe whether a household possessed a coupon that was not used. Therefore it is difficult to determine the cost of a counterfactual basket of goods.

[129]Chetty, Looney, and Kroft (2009) find that tax salience affects behavioural responses, with individuals responding more optimally (in terms of smaller utility losses) to sales-tax-inclusive prices compared to when sales taxes are added at the register. Therefore, interpretation of revealed preference results and tests for rationality depend on whether prices include tax or not.

[130]In other words, the aggregate price of category k in period t is defined as

$$p_t^k = \sum_{\forall i} \left(\frac{q_{i,t}^k}{\sum_{\forall i} q_{i,t}^k} \right) p_{i,t}^k.$$

[131]For example, if a product category contains several brands of varying qualities, then an increase in the unit value price could either mean that prices have increased or that consumers have switched to buying more expensive brands. Since it is likely that price and quality are correlated,

Tables 3.26-3.27 and Figure 3.18 summarise the data after aggregation (see Appendix 3A, Tables -3.31 for additional descriptive information). The aggregate prices across categories are quite similar in terms of mean and price range (possibly with the exception of pills), suggesting that a common cutoff would be plausible in this case[132]. Note that some categories (for example, pills and nuts) are purchased infrequently, so even after monthly aggregation, households generally do not purchase goods from every category in each month. Counterfactual prices for categories not purchased were calculated using the simple average of all prices in that category for that month. Thus, households who did not purchase from a particular category in a given month will face the same counterfactual price for that category.

Table 3.26: Observations per category (disaggregated)

Category	N
Bacon	4532
BBQ sauce	10573
Butter	11626
Cereal	15667
Chips/Crisps	18383
Coffee	5562
Crackers	9404
Eggs	10658
Ice cream	9632
Nuts	1799
Pills	1348
Pizza	4564
Sugar	4453
Total	108201

and that households will respond to price increases by switching to cheaper brands, then unit values will understate the original price increase.

[132] As mentioned in Section 3.1.1, the single cutoff assumption would be computationally convenient. For a given range, the GARP test only needs to be run once for each price in that range, for each household. If goods had different cutoffs, however, the number of computations increases exponentially with the number of goods and cutoffs, making a grid search for good-specific cutoffs infeasible in practice. For this 13-good dataset, searching for cutoffs in the range $0.01 - 1$ (99 values) would require 99^{13} computations per household, which is clearly infeasible without a supercomputer. For a robustness check, the revealed preference tests will be conducted with pills excluded from the list of goods.

Table 3.27: Household expenditure behaviour (disaggregated)

	Mean	SD	Min	Max
Mean # trips (per month)	4.14	1.70	1.50	17.54
Mean # categories (per month)	4.68	1.60	1.38	9.92
Mean # items (per trip)	6.20	3.08	1.53	17.90
Mean expenditure (per trip)	11.73	9.62	0	69.48

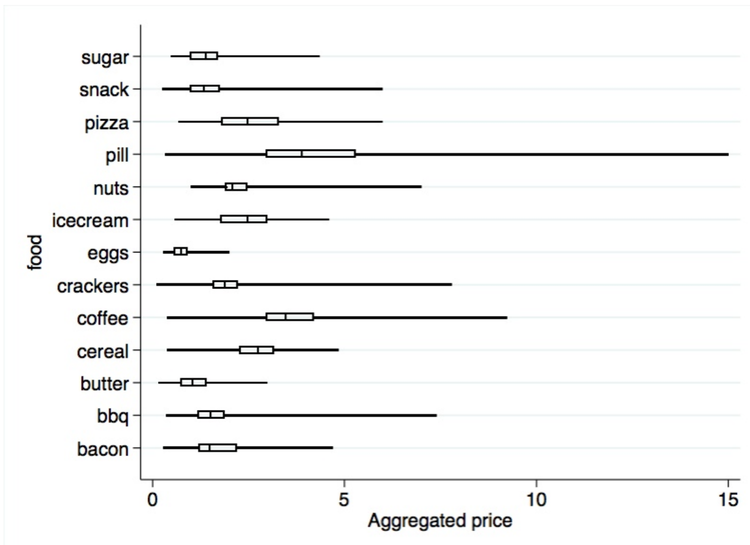

Figure 3.18: Aggregated prices (USD)

3.3.2 Revealed preference results

PI-GARP was tested for each household separately and for each price in the range $[0, 15.00]$, using 1-cent increments and assuming a common cutoff for all goods[133]. Table 3.28 shows the pass rates, power, and predictive success (PS) for both PI-GARP and standard GARP. Out of 494 households, 211 (43%) passed PI-GARP, while 135 (27%) passed standard GARP[134], which is equivalent to passing PI-GARP with a cutoff of 0. These results suggest that the

[133] The discussion of Figure 3.20 gives the reason why this interval was chosen.

[134] The pass rates for standard GARP differ from those of Echenique, Lee, and Shum (2012) because they define one time period as a set of four consecutive weeks (starting from the first week available) instead of actual calendar months, and they aggregate products according to brand name rather than according to product category. It is difficult to replicate the exact results of Echenique, Lee, and Shum (2012) from the raw data since they do not provide documentation on how they constructed their dataset (for example, exact details on aggregation and observations dropped from the raw data).

common cutoff assumption is reasonable in this setting[135]. Power is a measure of how easy it is to find data that fails the revealed preference test, and is considerably lower for PI-GARP (which places no restrictions on the cutoff) compared to GARP (which restricts the cutoff to 0). Predictive success indicates how 'good' a model is, according to the tradeoff between pass rate and power. Predictive success is negative for both models, meaning that they perform worse than a random number generator, and is larger (in absolute terms) for PI-GARP. In other words, PI-GARP has a higher pass rate than standard GARP because it is less restrictive, not because it makes more accurate predictions about observed behaviour. Although allowing for non-zero cutoffs within a large range ($\$0-\$15$) is much less restrictive than a zero cutoff, there is still some power to reject the model, so PI-GARP is falsifiable.

Table 3.28: Summary of revealed preference test results

	Pass	Power	PS
GARP	0.273	0.679	-0.048
	(0.446)	(0.141)	(0.445)
PI-GARP	0.427	0.150	-0.213
	(0.495)	(0.088)	(0.437)

Figure 3.19 presents the set and distribution of rationalisable cutoffs. The top panel of Figure 3.19 shows the set of rationalisable cutoffs for the first 50 (out of 211) households that passed PI-GARP, sorting according to household ID. The horizontal dotted lines show that the set of cutoffs for a given household can be non-convex, since the set of possible support prices is non-convex[136].

[135] It is possible that good-specific cutoffs could rationalise the remaining data, but it is too computationally burdensome to check, even for a small subset of cutoff values. To illustrate, with 13 goods and 2 given cutoffs, the number of permutations to check per household is already $2^{13} = 8192$, compared to 1 permutation assuming a common cutoff, so the computational burden increases exponentially with the number of cutoffs and goods.

As a robustness check, pills were dropped from the list of goods. The pass rates for GARP and PI-GARP with pills omitted were quite similar (0.267 and 0.433 respectively).

[136] This nonconvexity has two practical implications: 1. It motivates the focus on the minimum cutoff, because there may not be a maximum cutoff (if the data is rationalisable with a cutoff of $c = max(|\Delta p_t^k|)$ for $\forall t, k$, then it is trivially rationalisable by any larger cutoff), and 2. It would be computationally burdensome to find the set of rationalisable cutoffs if cutoffs were good-specific, because it does not suffice to find the minimum and maximum cutoff.

Although the size of the set varies greatly across households, the bottom panel indicates that most rationalisable cutoffs are concentrated in the smaller values ($\$0 - \$0.5$), and the minimum cutoff needed to rationalise the data also falls within this range. Also, the CDF of the minimum cutoff (Figure 3.20) suggests that it is not possible to rationalise any data simply by increasing the size of the cutoff: if the data is not rationalisable by a plausible common cutoff, then it is likely that it is not rationalisable at all. As discussed in Section 3.2, any cutoff greater than or equal to the maximum price change across all goods would generate the same support prices, so any data that cannot be rationalised by a cutoff in the range $[0, max_{t,k}\{\triangle p_t^k\}]$, then it cannot be rationalised at all. For this particular data, since the maximum price difference between any two consecutive periods is less than \$15.00, we know for certain that the remaining households will not pass PI-GARP for any higher common cutoff. These are both plausible implications of the price-inattention model.

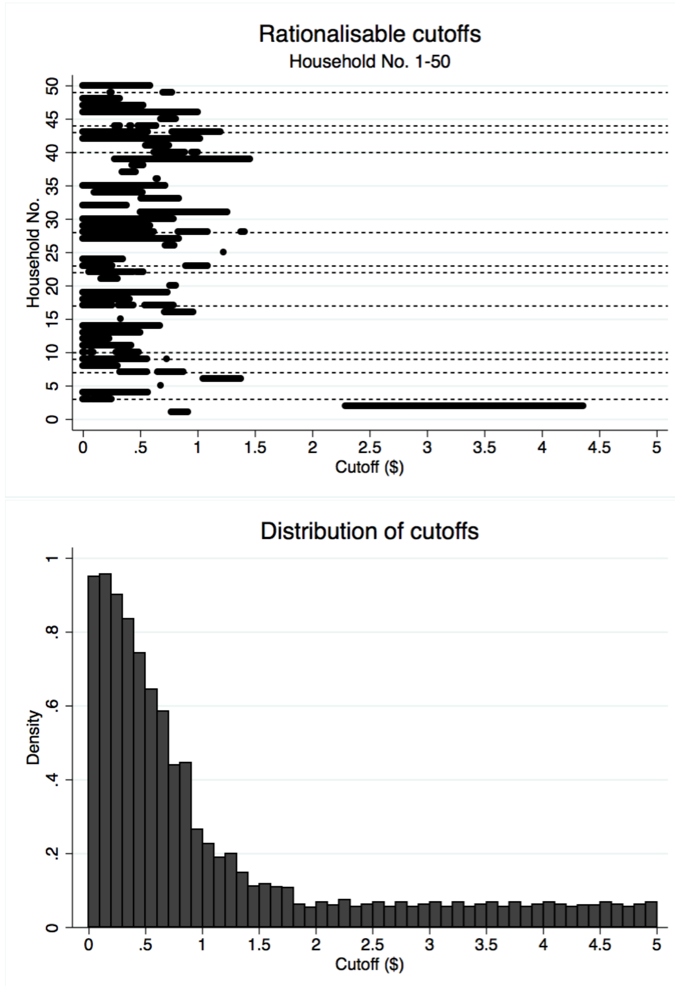

Figure 3.19: Rationalisable cutoffs: examples and frequency distribution

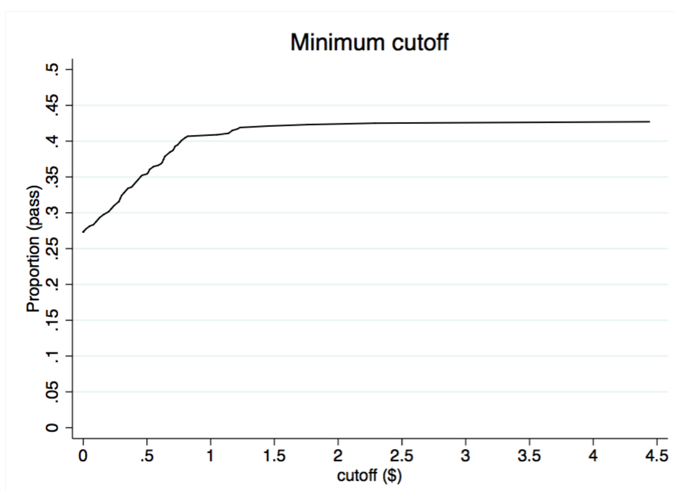

Mean	SD	Min	Max	N
0.21	0.46	0	4.44	211

Figure 3.20: Minimum cutoff: CDF and summary statistics

3.3.3 Rationalisability and observables

Does price inattention or the minimum cutoff depend on household characteristics? Table 3.29 shows the results of regressions of the pass rate and minimum cutoff on available demographic variables and information about purchasing behaviour (standard errors in brackets, and coefficients with $p \leq 0.1$ in bold font). Having a larger family significantly increases the likelihood of passing, possibly because larger households have to be more economical with spending and are forced to think rationally about every purchase. Surprisingly, being more educated significantly decreases the likelihood of passing.

The results for the regression on minimum cutoff are more intuitive. Having a larger family decreases the minimum cutoff, for the same reasons mentioned above. Being richer significantly increases the cutoff, presumably because higher-income households can afford not to pay too much attention to the actual prices. Similarly, buying more items on average in one trip, or buying a wider variety of categories per trip is associated with higher cutoffs, since households are less likely to notice price changes of any one item or category in their basket. Lower minimum cutoffs are associated with being more educated and shopping more frequently, since these households are more likely to notice price changes.

Table 3.29: Correlation between rationality, the minimum cutoff, and household characteristics

	Pass (0/1)		Min cutoff
family size	**0.087**	family size	-0.001
	(0.050)		(0.025)
income	0.009	income	**0.023**
	(0.027)		**(0.012)**
black	0.261	black	-0.049
	(0.173)		(0.065)
education	**-0.188**	education	-0.057
	(0.101)		(0.047)
mean # trips	-0.053	mean # trips	-0.014
	(0.040)		(0.020)
mean # categories	-0.055	mean # categories	0.021
	(0.053)		(0.027)
mean # items	0.037	mean # items	0.002
	(0.027)		(0.011)
constant	0.293	constant	0.201
	(0.326)		(0.128)
Pseudo-R^2	0.026	R^2	0.026

3.3.4 Perceived cost-of-living indices

The presence of price inattention requires modifying the prices used for tests of rationality, as well as using perceived prices instead of actual prices to conduct welfare analysis for price-inattentive individuals or groups of consumers that are partially composed of price-inattentive individuals. Since prices are essential for calculating statistics such as the cost of living, or measuring the compensating and equivalent variation, acting on misperceived prices could imply very different estimates of welfare changes, especially if the magnitude of price changes is small (within the individual's cutoff). While conducting welfare analysis is beyond the scope of this chapter, this section explores the implications of price-inattentive behaviour on perceived cost-of-living indices.

Cost-of-living indices compare the minimum cost required to maintain a given level of utility under one set of prices, relative to the cost under another set of prices (base-period prices,

$\mathbf{p}_B$). Using $C(\mathbf{p}_t, u)$ to denote the minimum cost of achieving utility level u given prices at time t, the 'true' cost of living index is defined as the ratio of the cost functions in both periods (Konus, 1924)[137]:

$$P_{K,t} = \frac{C(\mathbf{p}_t,u)}{C(\mathbf{p}_B,u)} \text{ for } t \geq 0$$

The cost function $C(\mathbf{p}_t, u)$ is unobserved but can be approximated using observed prices and quantities, assuming that consumers are neoclassical utility maximisers, so that $\mathbf{q}_B$ and $\mathbf{q}_t$ are both cost-minimising bundles[138]. The Laspeyres price index (P_L) compares the cost of purchasing the base-period bundle ($\mathbf{q}_B$) under base-period and period-t prices, while the Paasche price index (P_P) compares the cost of purchasing the period-t bundle ($\mathbf{q}_t$) under base-period and period-t prices:

$$P_{L,t} = \frac{\mathbf{p}_t'\mathbf{q}_B}{\mathbf{p}_B'\mathbf{q}_B} \text{ for } t \geq 0$$

$$P_{P,t} = \frac{\mathbf{p}_t'\mathbf{q}_t}{\mathbf{p}_B'\mathbf{q}_t} \text{ for } t \geq 0$$

The Laspeyres price index bounds the true cost-of-living index from above, since any bundle that is optimal at prices $\mathbf{p}_t$ must cost weakly less than $\mathbf{p}_t'\mathbf{q}_B$, and likewise the Paasche price index provides a lower bound on the true cost-of-living index, since any bundle that is optimal at prices $\mathbf{p}_B$ must cost weakly less than $\mathbf{p}_B'\mathbf{q}_t$.

The Lasperes and Paasche price indices are sensitive to the choice of base period and thus fail the 'time reversal' property, meaning that the value of the index depends on the order in which the price changes are made[139]. Two standard indices that overcome this issue are the Fisher (P_F) and Tornqvist (P_T) indices, which are defined below:

$$P_{F,t} = \sqrt{P_{L,t}P_{P,t}} \text{ for } t \geq 0$$

$$P_{T,t} = \prod_{i=1}^{K} (\frac{p_{i,t}}{p_{i,B}})^{\frac{1}{2}\{w_{i,B}+w_{i,t}\}} \text{ for } t \geq 0$$

[137]This cost of living index is directly related to compensating and equivalent variation: $CV = C(p_t, u_0) - C(p_0, u_0)$ and $EV = C(p_t, u_t) - C(p_0, u_t)$.

[138]That is, they are both solutions to the cost-minimisation problem $\min_{\mathbf{q}}\{\mathbf{p}'\mathbf{q} : U(\mathbf{q}) \geq u\}$ at the given prices.

[139]For example, if the prices and quantities in time 0 and time t were interchanged, then the resulting Laspeyres price index ($P_{L,0} = \frac{\mathbf{p}_0'\mathbf{q}_t}{\mathbf{p}_t'\mathbf{q}_t}$) would not be the reciprocal of the original price index, so the ordering of time periods matters.

where $w_{i,t}$ is the expenditure share of good i in period t ($w_{i,t} = \frac{p_{i,t}q_{it}}{\sum_{i=1}^{K} p_{it}q_{it}}$).

Indices may be chained or non-chained, depending on the choice of base period. Non-chained indices usually take the first period to be the base period ($\mathbf{p}_B = \mathbf{p}_0$), whereas chained indices take the previous period to be the base period ($\mathbf{p}_B = \mathbf{p}_{t-1}$). Chained indices are useful for comparing period-to-period changes, and price changes over time when there are large shifts in prices and quantities compared to the initial period. Also, as with the price-inattentive consumer, chained indices only use information in the current and previous period.

Cost-of-living indices depend on the assumption that a consumer's preferences can be represented by a well-behaved utility function, and that they behave rationally according to these preferences. If consumers do not behave rationally according to standard economic theory, for example, by acting on perceived rather than actual prices, then these indices need to be modified in order to conduct welfare analysis, calculate perceived cost-of-living indices, and infer consumers' perceptions of inflation.

When price changes are small, price-inattentive households may not notice the actual prices and hence will also misperceive their cost of living. The price-inattentive household evaluates their cost of living according to perceived prices ($\tilde{\mathbf{p}}_t$) instead of actual prices ($\mathbf{p}_t$), and behaves as if they solve $\min_{\mathbf{q}}\{\tilde{\mathbf{p}}'\mathbf{q} : U(\mathbf{q}) \geq u\}$ (defined as $C(\tilde{\mathbf{p}}, u)$), rather than $\min_{\mathbf{q}}\{\mathbf{p}'\mathbf{q} : U(\mathbf{q}) \geq u\}$ ($C(\mathbf{p}, u)$). The price-inattentive equivalents of the four price indices are therefore obtained by replacing actual prices with perceived prices:

$$\tilde{P}_{L,t} = \frac{\tilde{\mathbf{p}}_t'\mathbf{q}_B}{\tilde{\mathbf{p}}_B'\mathbf{q}_B} \text{ for } t \geq 2$$

$$\tilde{P}_{P,t} = \frac{\tilde{\mathbf{p}}_t'\mathbf{q}_t}{\tilde{\mathbf{p}}_B'\mathbf{q}_t} \text{ for } t \geq 2$$

$$\tilde{P}_{F,t} = \sqrt{\tilde{P}_{L,t}\tilde{P}_{P,t}} \text{ for } t \geq 2$$

$$\tilde{P}_{T,t} = \prod_{i=1}^{K} \left(\frac{\tilde{p}_{i,t}}{\tilde{p}_{i,B}}\right)^{\frac{1}{2}\{\tilde{w}_{i,B}+\tilde{w}_{i,t}\}} \text{ for } t \geq 2$$

where the second observed period is now the base period ($\mathbf{p}_B = \mathbf{p}_2$) for non-chained indices. One way to see this is that under price-inattentive behaviour, we observe a combination of actual budget constraints (either because all the prices were large enough to be noticed, or the consumer re-optimised after overspending) and underspends (because not all price changes were noticed). To obtain a consumer's perceived budget constraint and calculate their perceived cost of living, we adjust the budget constraints according to perceived prices instead of using observed prices.

Given the minimum cutoff for a household, we can construct the set of perceived prices and hence the perceived price indices for each household. Figure 3.21 compares the perceived ('fuzzy') index with the actual ('standard') index for some households, using the Laspeyres non-chained index as an example. If the household noticed all price changes, then $\tilde{P}_{L,t} = P_{L,t}$. Conversely, for the first household (left panel), the minimum cutoff is larger than all the price changes, so the household never notices any price changes. In this case, $\tilde{P}_{L,t}$ equals the lagged actual index ($P_{L,t-1}$), because the household only learns about the actual prices at the end of the period. The second household (right panel) notices some price changes, but not all. In this case, assuming that the household correctly perceives prices in the base period, $\tilde{P}_{L,t}$ is a combination of $P_{L,t-1}$ in periods following a large enough price change, but otherwise can be completely different in periods following a small price change. If the household incorrectly perceives prices in the base period, the index value may be different in all periods, even if they correctly perceive prices in a given period.

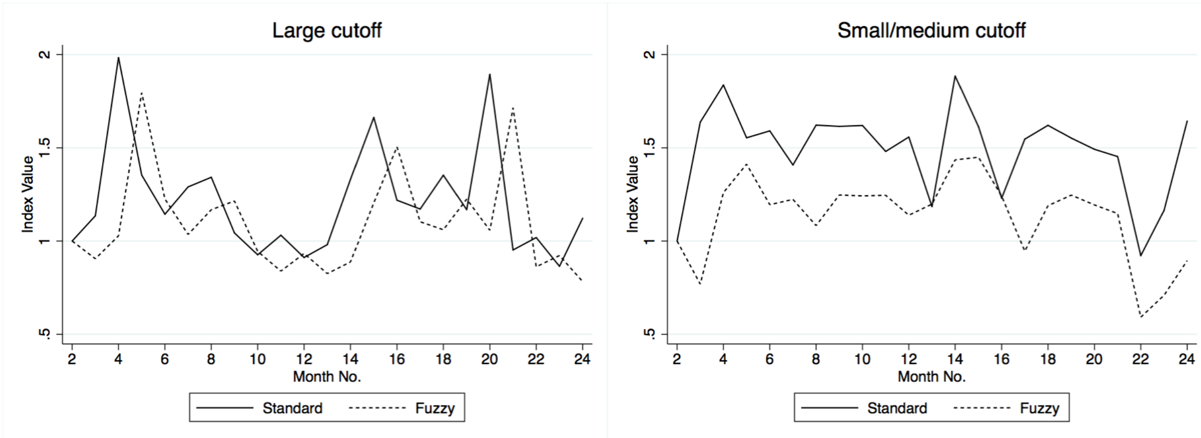

Figure 3.21: Examples of household-level Laspeyres indices

Figures 3.27 and 3.28 (Appendix 3A) show 2D contour plots of the distribution of values for each index (Figures 3.29 and 3.30 (Appendix 3A) show 3D surface plots of the same distributions for a more interesting visualisation). In any given month, there is a considerable range across households for all indices, though most values are concentrated around the mean value in that month.

In an economy made of price-inattentive and neoclassical households, what effect does the former have on the aggregate cost-of-living index? Figure 3.22 shows the four price indices (solid lines, shaded 95% confidence interval) and their perceived counterparts (dashed lines, dotted 95% confidence interval), averaged across all households that passed PI-GARP (either a simple average or a weighted average according to the expenditure in that month). Both the unweighted and weighted perceived indices closely track their standard counterparts, with departures occurring when prices do not change much. These departures are larger for the weighted indices, suggesting that the price-inattentive households spent more than neoclassical households.

Figure 3.23 shows the chained version of these indices, which have a 'short memory' property similar to the price-inattentive household's behaviour (only using information from the current and previous period). Since the base period is the previous period, the perceived index will almost coincide with the standard index if the price changes in two consecutive periods are large enough to be noticed by most households. As with the non-chained indices, the chained perceived and standard indices generally do not differ by much, unless the standard index does not change by much, which is likely because most households that pass PI-GARP have a small cutoff (less than $1). Unlike in Figure 3.22, these departures are larger for the unweighted indices.

The link between good-specific price inattention and deviations from the standard index value is shown in Figure 3.24. Differences in the household-level index value are plotted against the net number of noticed good-specific price changes, calculated as the number of noticed price increases noticed minus the number of noticed price decreases[140]. A price change is noticed only if it is above the household-specific cutoff, which is taken to be the minimum rationalisable cutoff. As expected, deviations from the standard index value are smaller and near zero when prices of most goods either increase or decrease and are noticed. For the non-chained indices, the perceived index value exactly equals the standard index value if all price changes are noticed in period 2 and the current period, whereas for chained indices, this occurs only if all price changes are noticed in 2 consecutive periods, and hence the deviations are less concentrated around zero even when most price changes are noticed.

Figures 3.25 and 3.26 show the annual inflation rate, calculated from the aggregate cost-of-living indices in Figures 3.22 and 3.23. The perceived inflation rates (dashed lines) deviate from the actual inflation rates (solid lines) when inflation is low, and can differ by a percentage point or more, but closely track the actual inflation rate when inflation is high.

Overall, these results indicate that consumers tend not to notice inflation and generally behave in a price-inattentive manner when it is low (such that most micro-level price changes go unnoticed), but have more accurate price perceptions and generally behave neoclassically

[140] In this 13-good example, this measure can potentially range from -13 to 13, although as noted in Section 3.3.1, most households did not buy all 13 goods in each period.

when inflation is high. During periods of low inflation, the perceived costs of living can differ greatly from the actual cost of living, so welfare analysis and evaluations of changes in living standards may also be substantially different.

3.3.5 Extensions

There are some natural extensions to the model developed in this chapter. Firstly, the process by which the cutoff is determined can be modelled explicitly, for example, as the optimal level of inattention given the utility cost of inattention and the information processing cost of paying attention (as in Gabaix (2014)), or as the result of rounding product-level price changes to the nearest round number. It is possible that doing so would place restrictions on the range of possible cutoffs that can rationalise the data, and hence reduce the computational burden of the PI-GARP test.

Secondly, it would be worth investigating whether product promotions or changes in household characteristics affect the cutoff in either the short or long run. For example, discounts in specific products may make consumers more attentive to prices of that product in the future, but may also make consumers more attentive of the prices of other products during that shopping trip. These effects would have welfare implications for product marketing. This extension would be possible if there were complete information on coupons and discounts, or panel data on household characteristics.

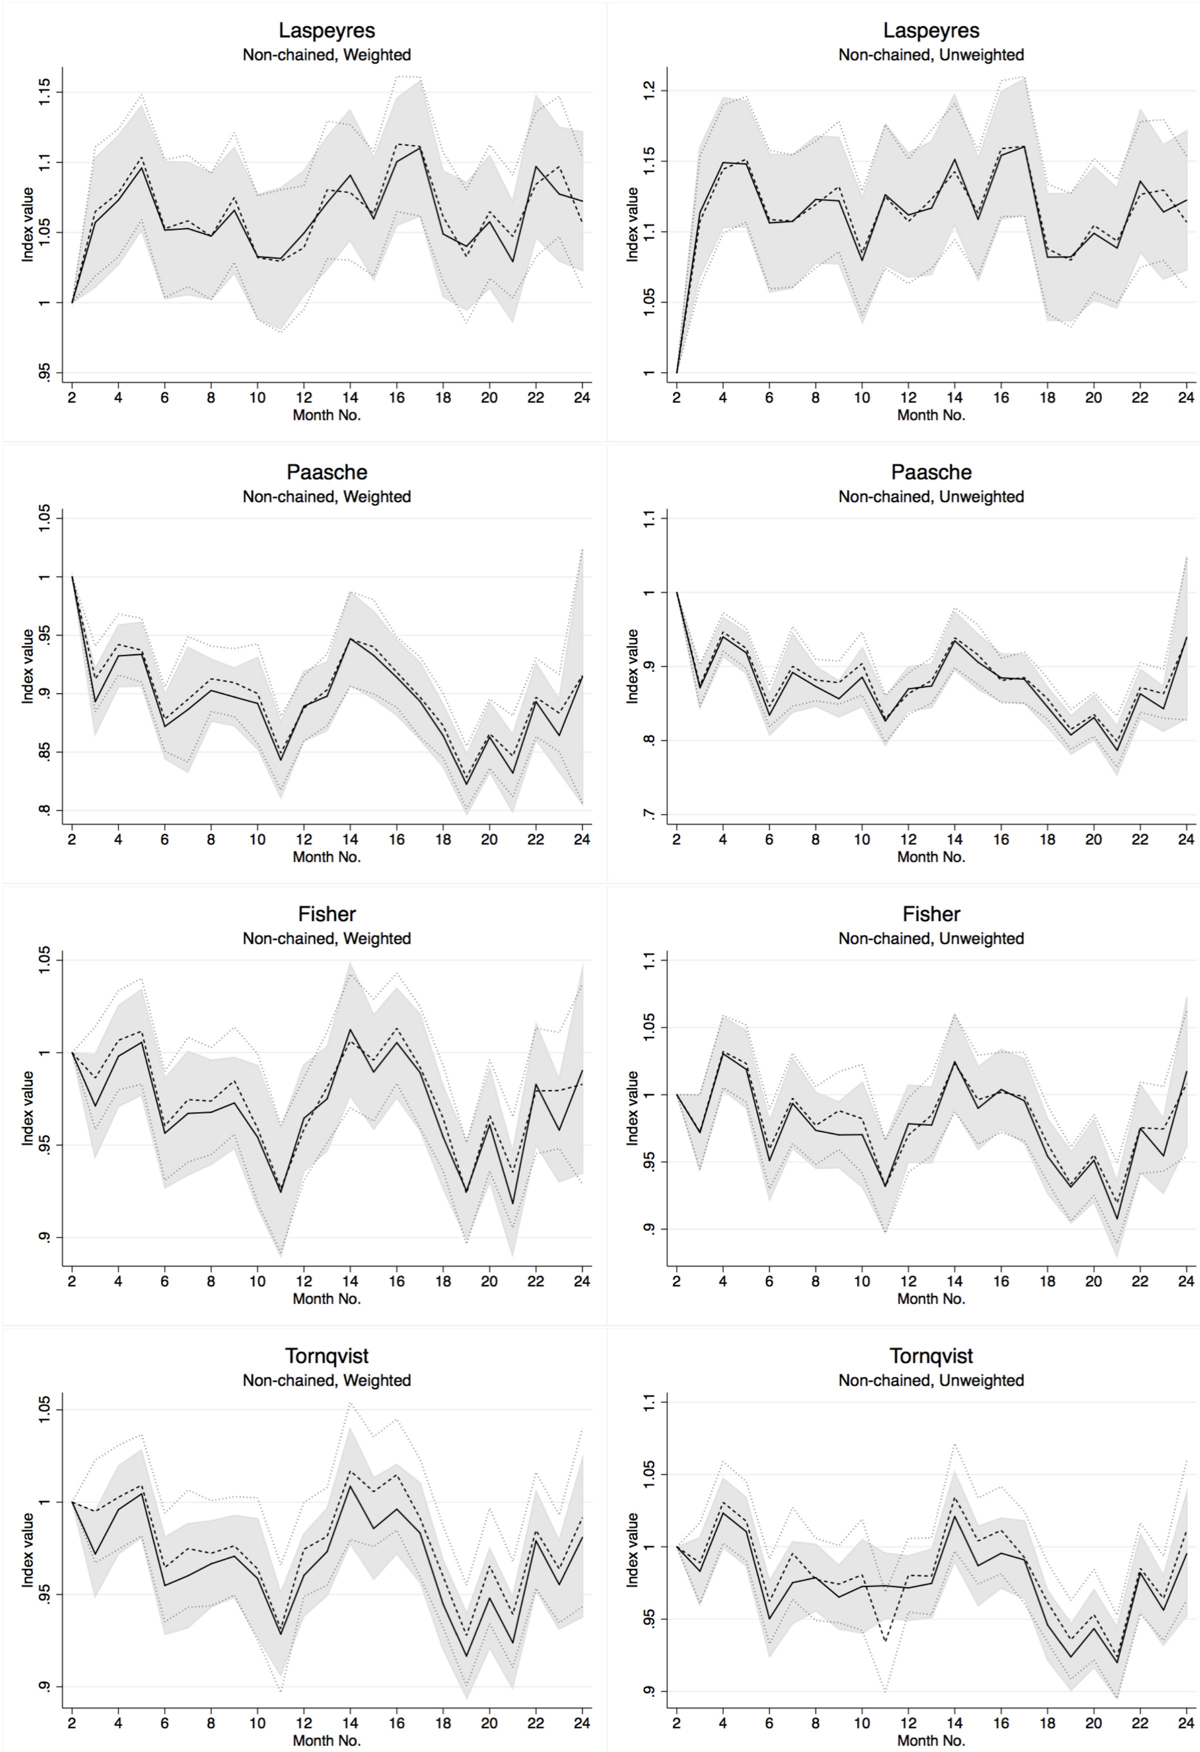

Figure 3.22: Standard and perceived price indices (Non-chained)

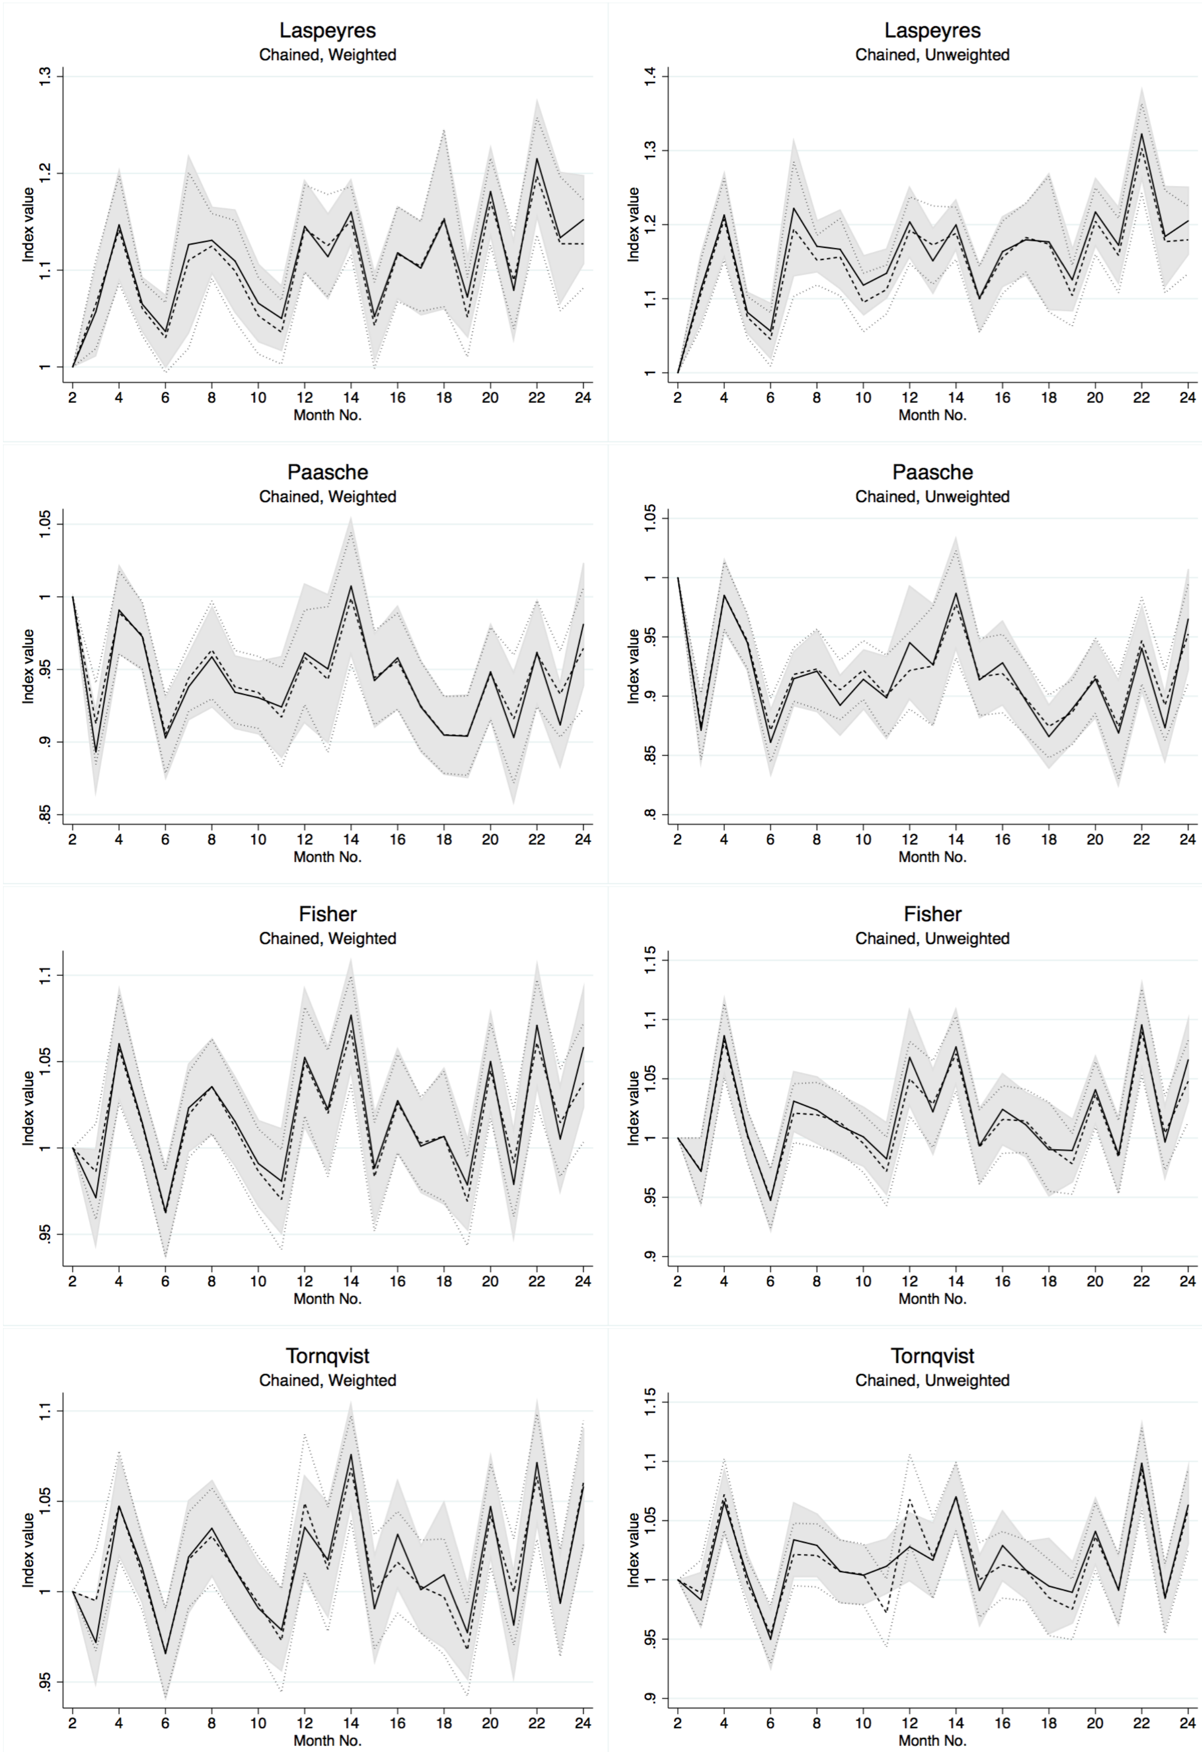

Figure 3.23: Standard and perceived price indices (Chained)

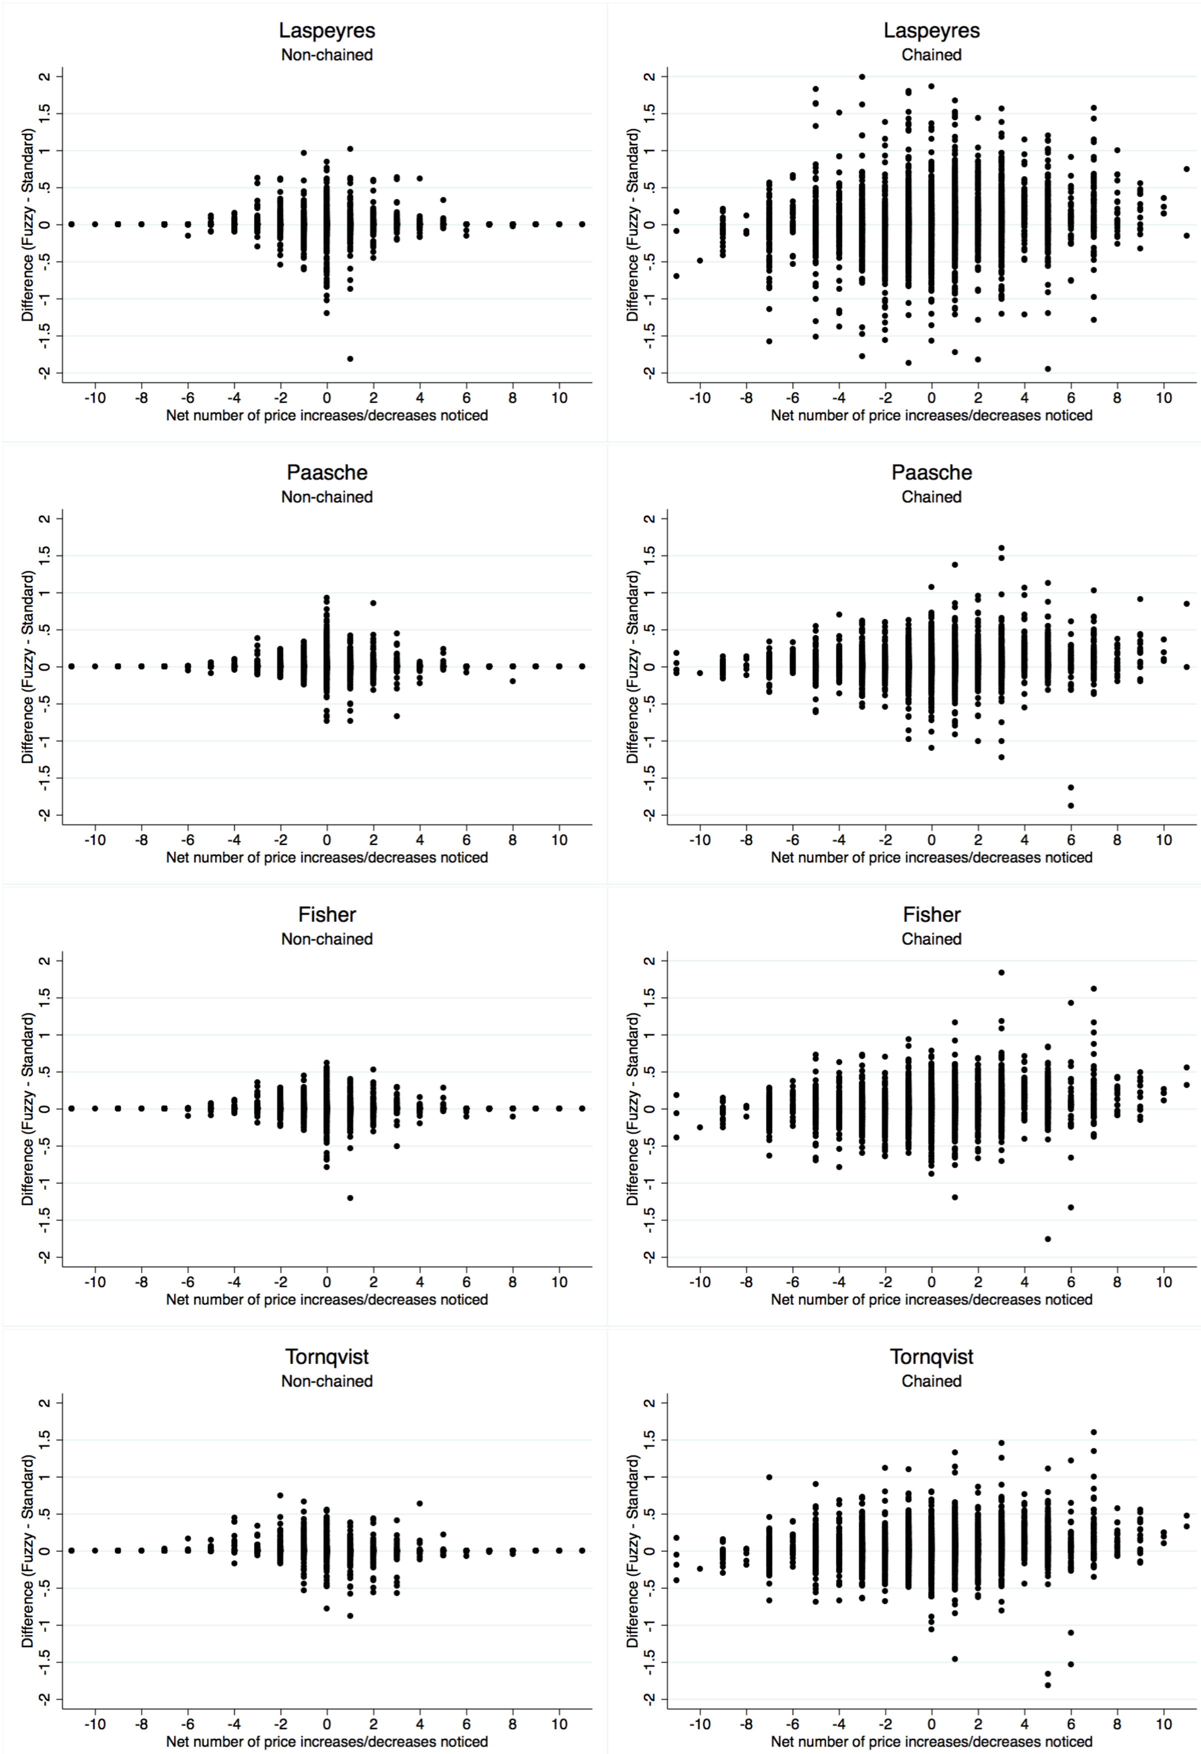

Figure 3.24: Differences between standard and perceived indices, depending on number of noticed price changes

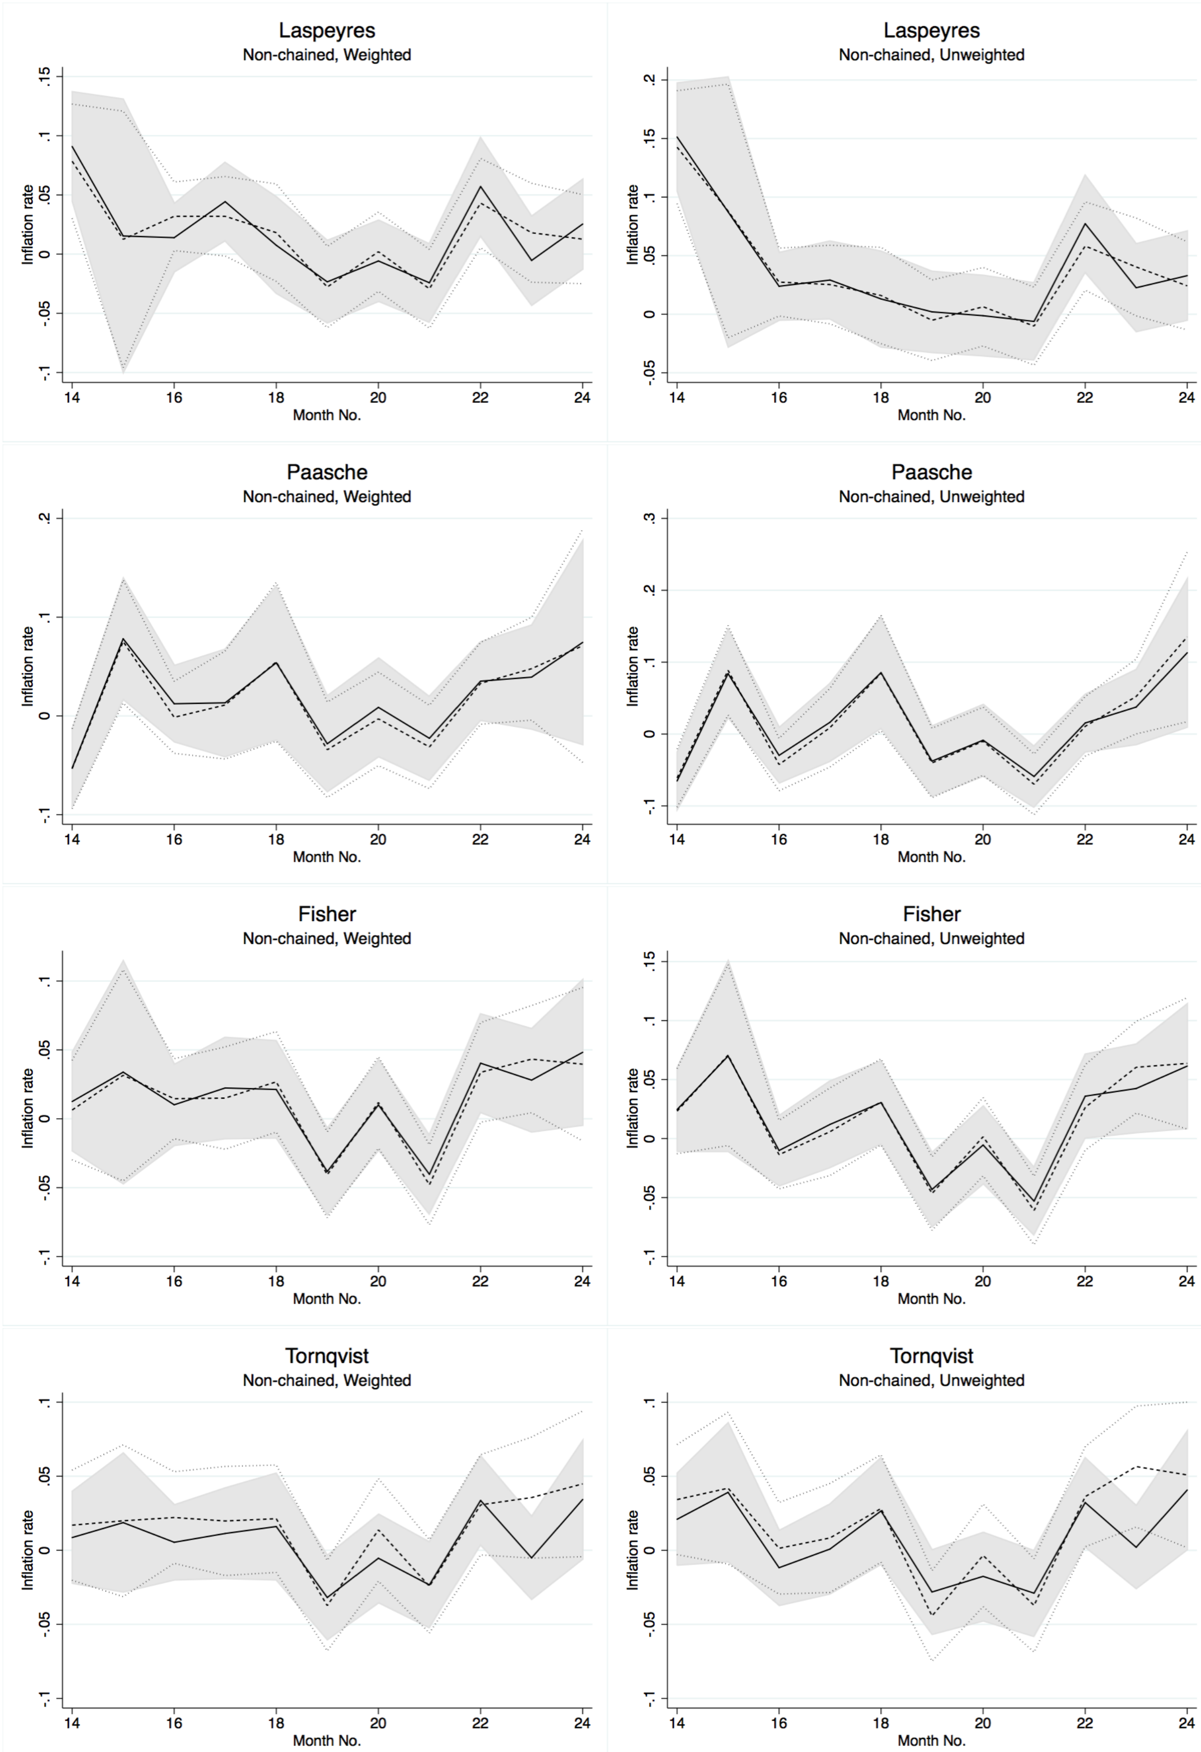

Figure 3.25: Annual inflation rates (Non-chained)

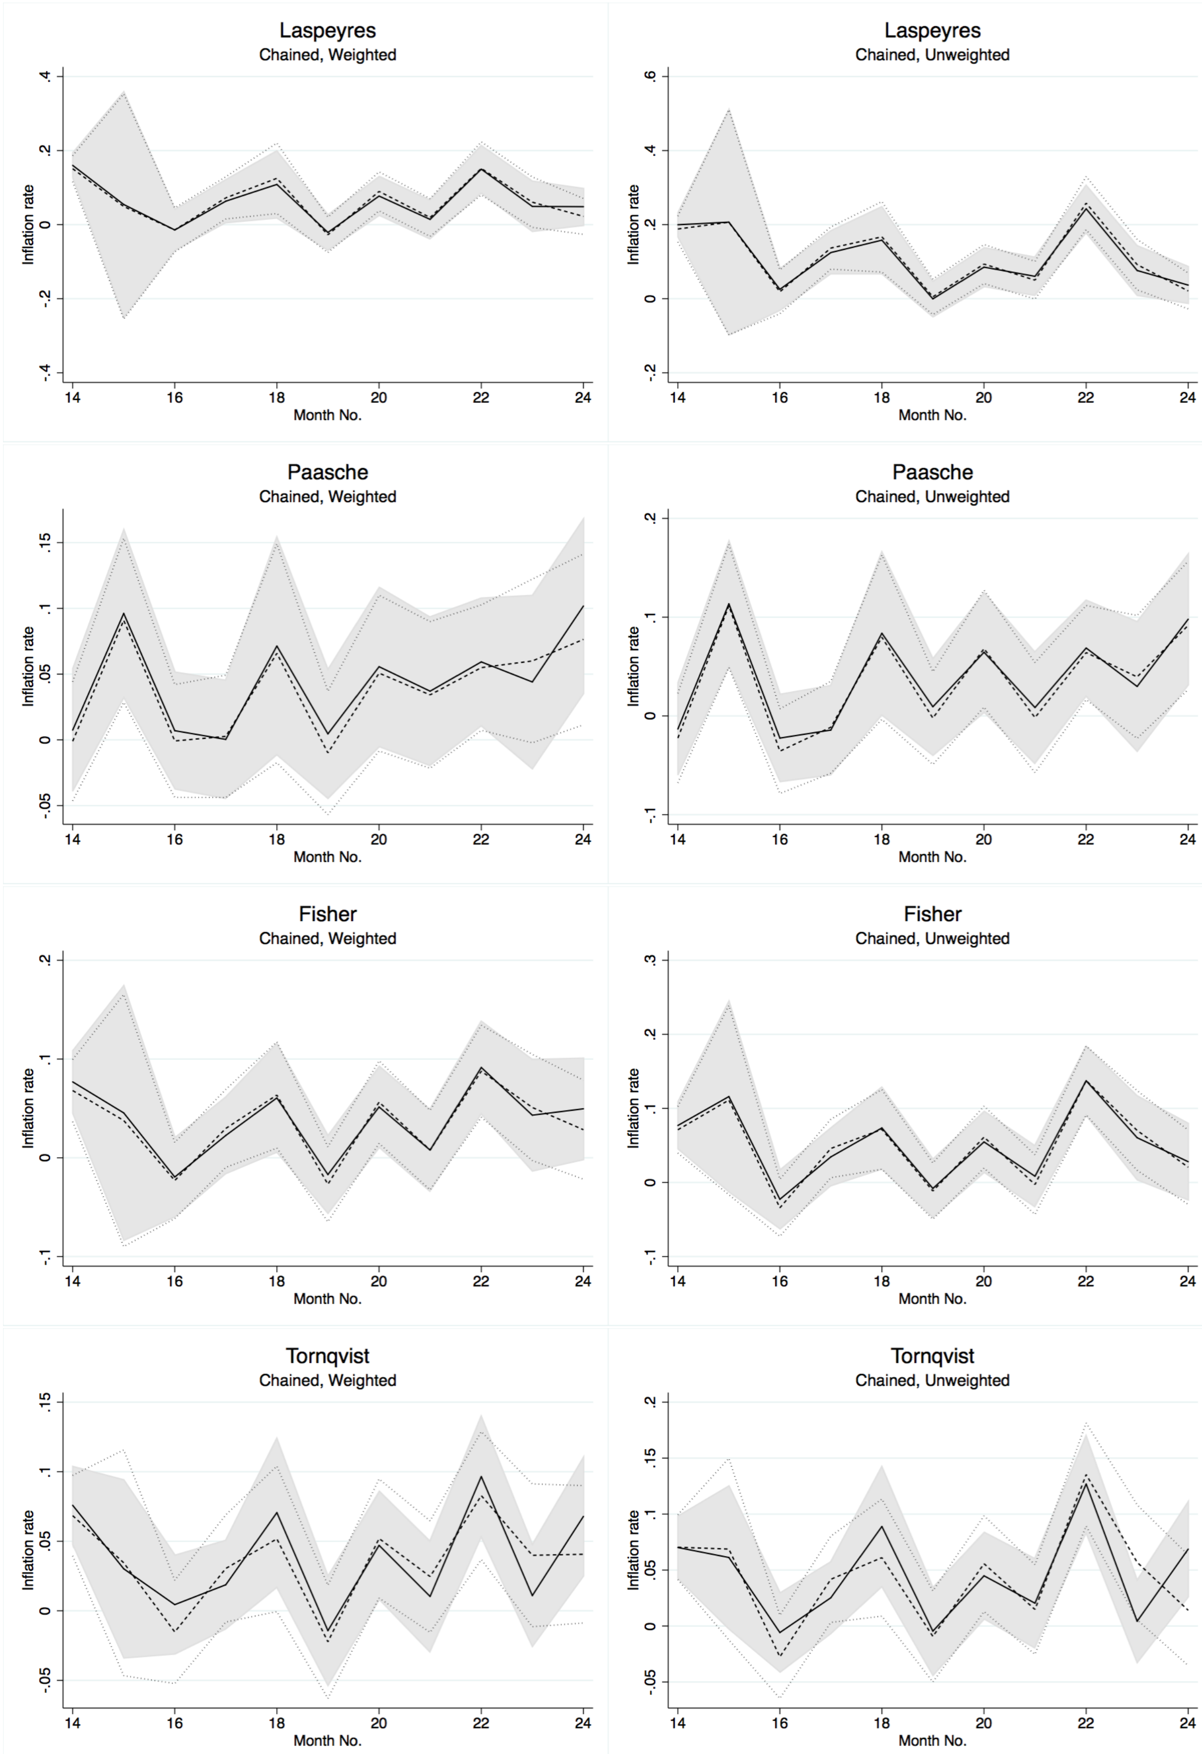

Figure 3.26: Annual inflation rate (Chained)

3.4 Conclusion

This chapter explores the effect of price inattention on the perception of costs of living and annual inflation rates, using a behavioural model in which consumers only notice price changes if they are 'large enough' and may not notice 'small' price changes. This price-inattention model is characterised by short-term memory, an asymmetric response to overspending and underspending, and a price cutoff, which can be good-specific or common to all goods. These features capture the notions of anchoring, partial adjustment towards the truth, and framing effects, which previous studies have shown to be part of behavioural decision-making. Necessary and sufficient conditions for data to be consistent with this price-inattention model are derived, which are easy to test given reasonable restrictions on the parameters.

A data application, using the Stanford Basket Dataset, demonstrates that the price-inattention model has some desirable properties: (1) it is plausible, meaning that the minimum cutoff found for households falls within a reasonable range; and (2) it is falsifiable, so not all behaviour can be rationalised with a large enough cutoff. By comparing perceived cost-of-living indices and annual inflation rates with their standard counterparts, this price-inattention model also generates the intuitive result that consumers have more accurate perceptions of inflation and generally behave neoclassically when inflation is high, but may not perceive inflation accurately when it is low. These results have important implications for conducting welfare analysis when consumers are not fully attentive to price changes.

Appendix 3A

Table 3.30: Supplementary summary tables: Household demographics

Race	N
No response	1
White	418
Black	66
Other	9
Total	494

Education	N
No response	2
Some grade school	8
Completed grade school	68
Completed high school	304
Completed college	112
Total	494

Household size	N
1	59
2	177
3	109
4	83
5	40
6+	26
Total	494

Household income (USD)	N
No response	6
$<$ 10000	35
10000-11999	26
12000-14999	30
15000-19999	53
20000-24999	48
25000-34999	76
35000-44999	79
45000-54999	63
55000-64999	25
65000-74999	23
75000+	30
Total	494

Table 3.31: Aggregated prices

Category	Mean	SD	Min	Max
Bacon	2.19	0.47	0.29	4.69
BBQ sauce	1.72	0.39	0.36	7.39
Butter	1.21	0.35	0.15	2.99
Cereal	2.84	0.52	0.39	4.85
Chips/Crisps	1.49	0.43	0.25	5.99
Coffee	3.73	0.71	0.39	9.22
Crackers	1.94	0.37	0.1	7.79
Eggs	0.84	0.19	0.29	1.99
Ice cream	2.79	0.63	0.59	4.59
Nuts	2.83	0.35	0.99	6.99
Pills	4.70	0.68	0.33	14.99
Pizza	2.83	0.49	0.69	5.99
Sugar	1.29	0.24	0.49	4.34

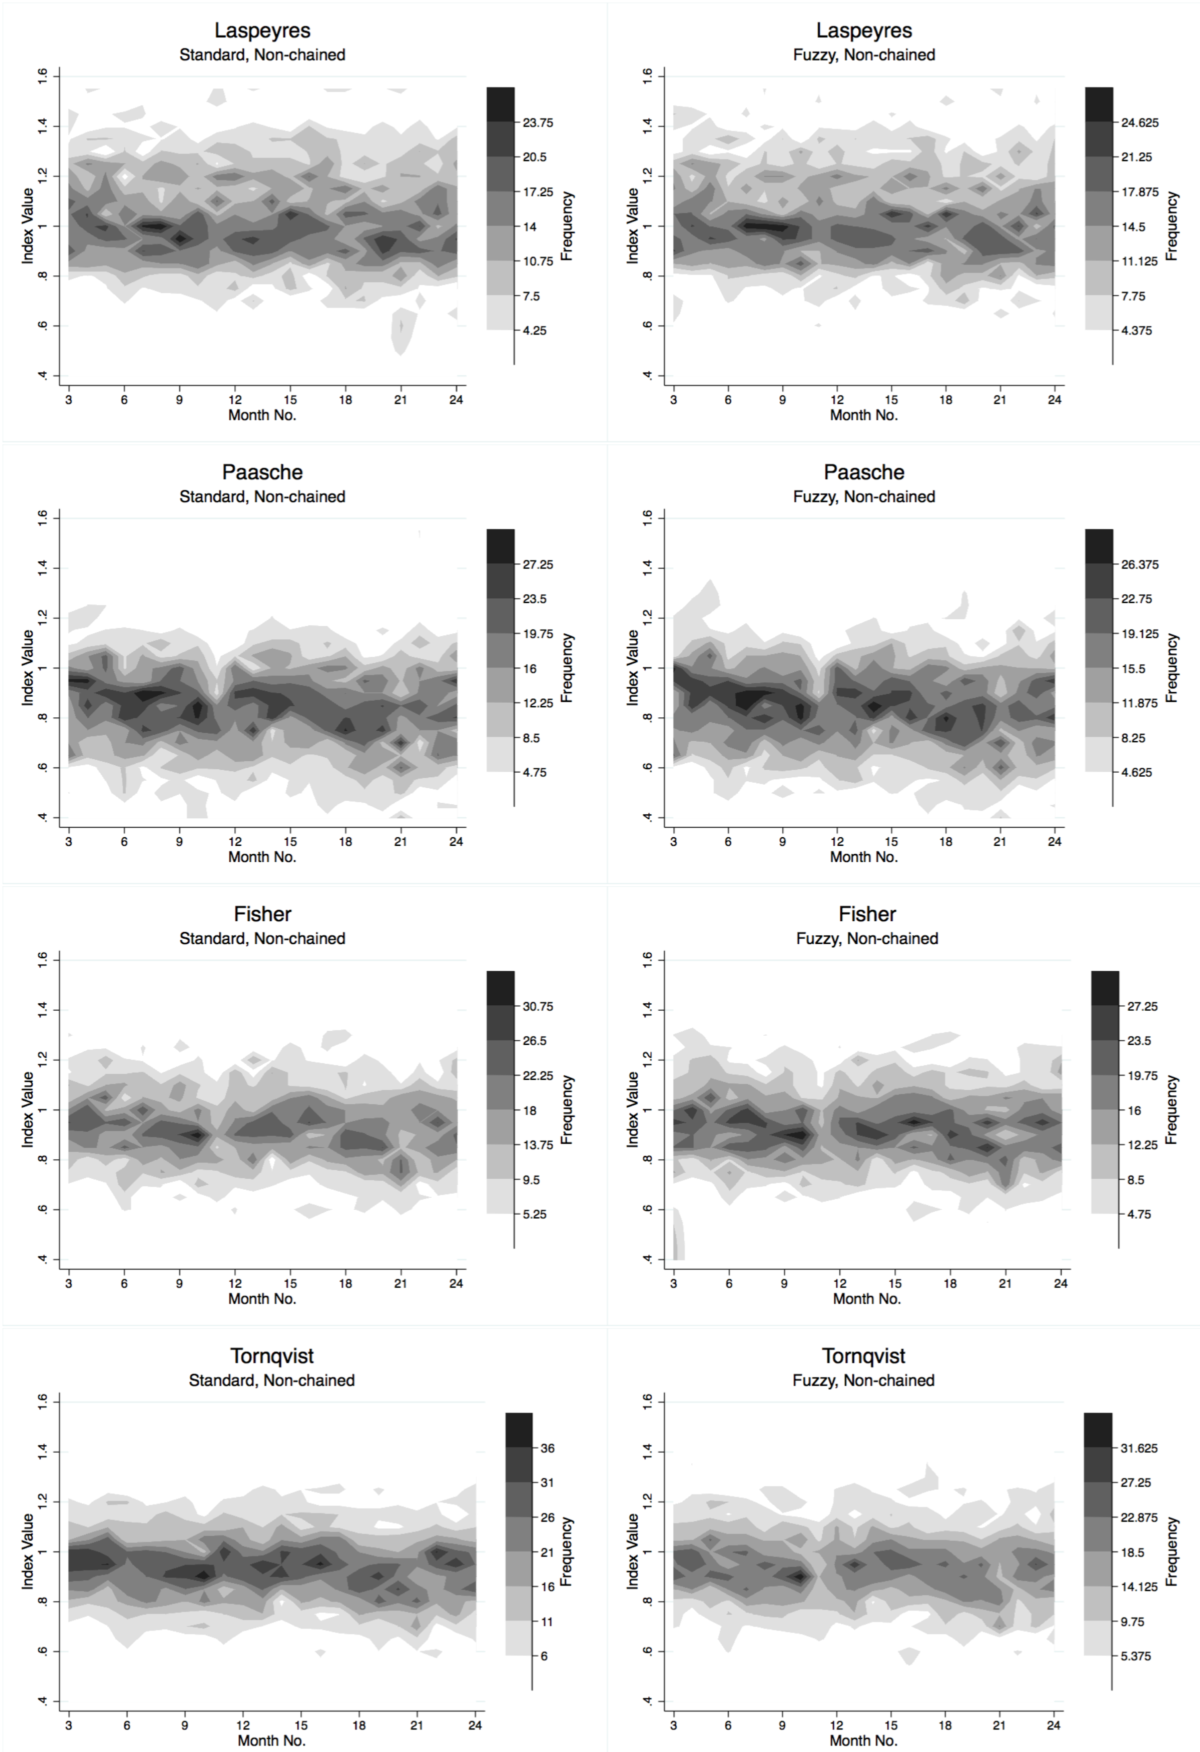

Figure 3.27: 2D contour plots (Non-chained)

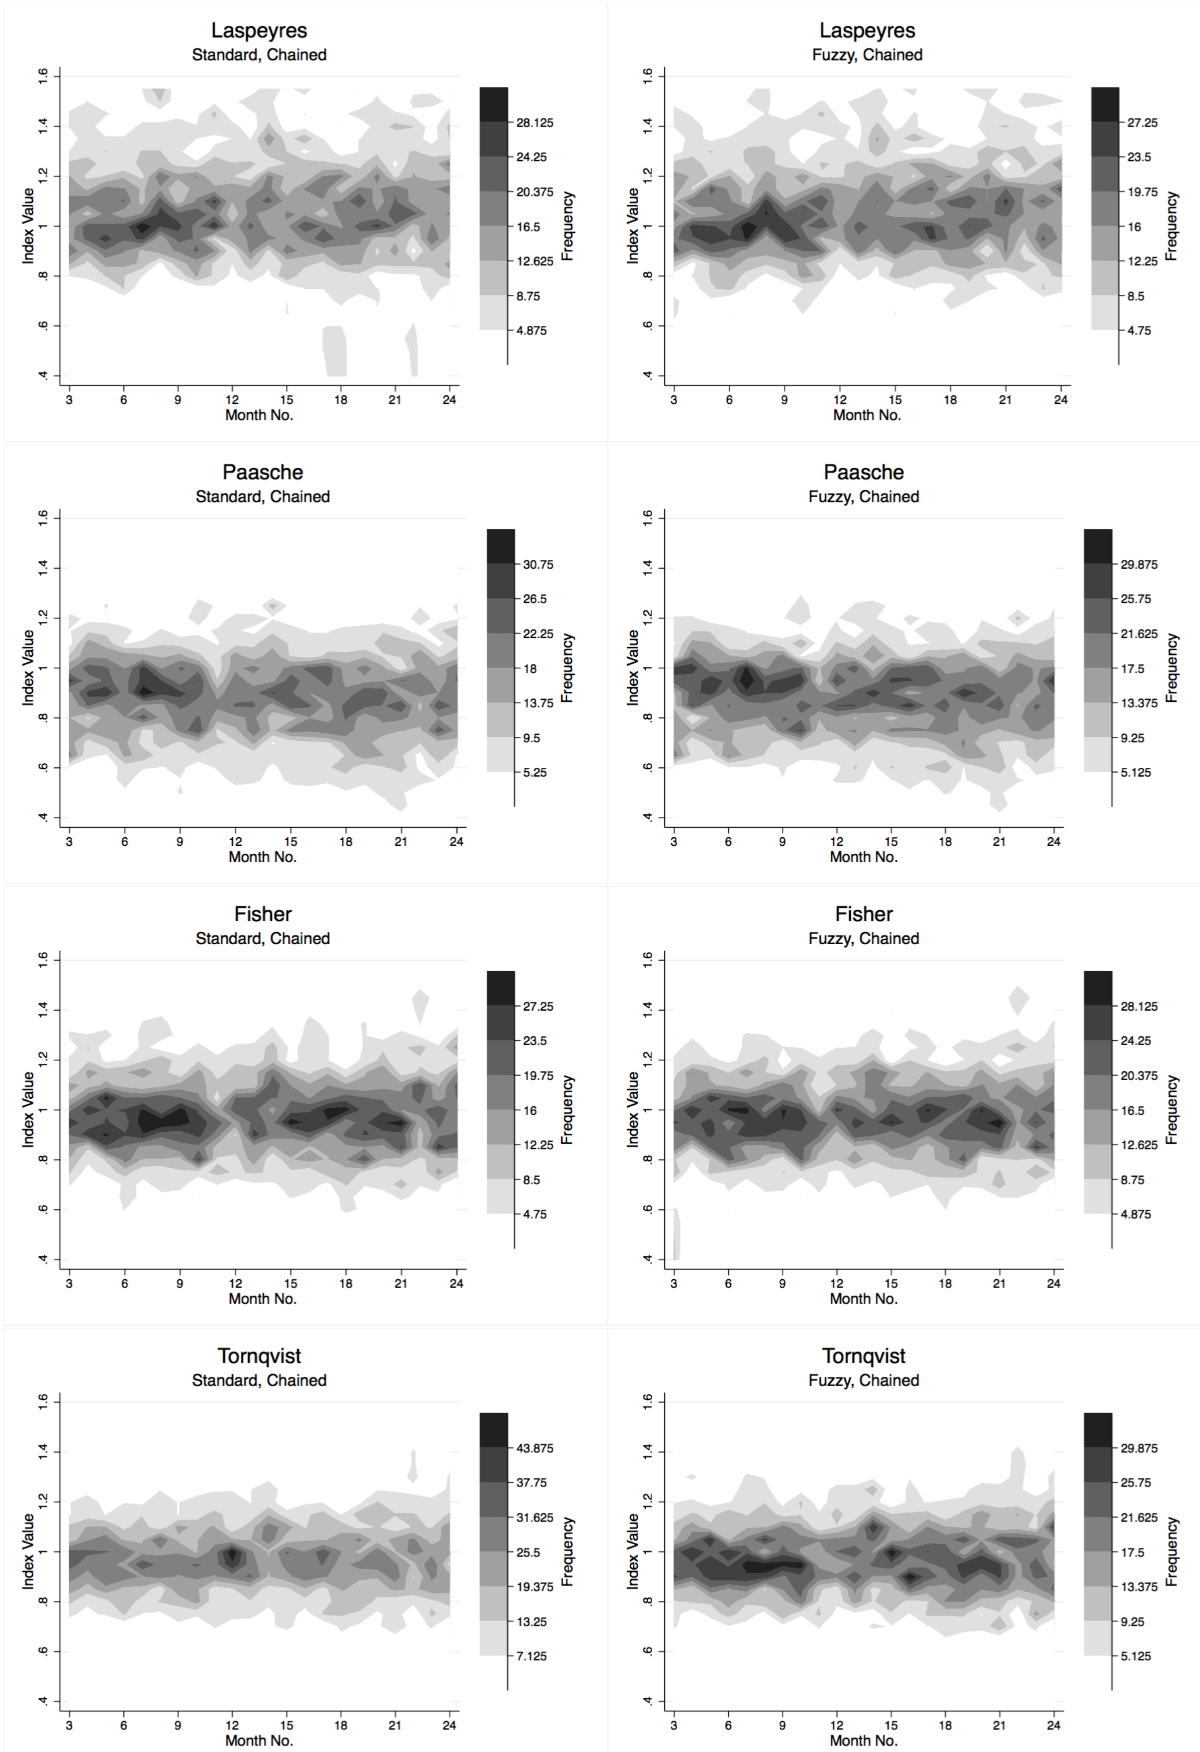

Figure 3.28: 2D contour plots (Chained)

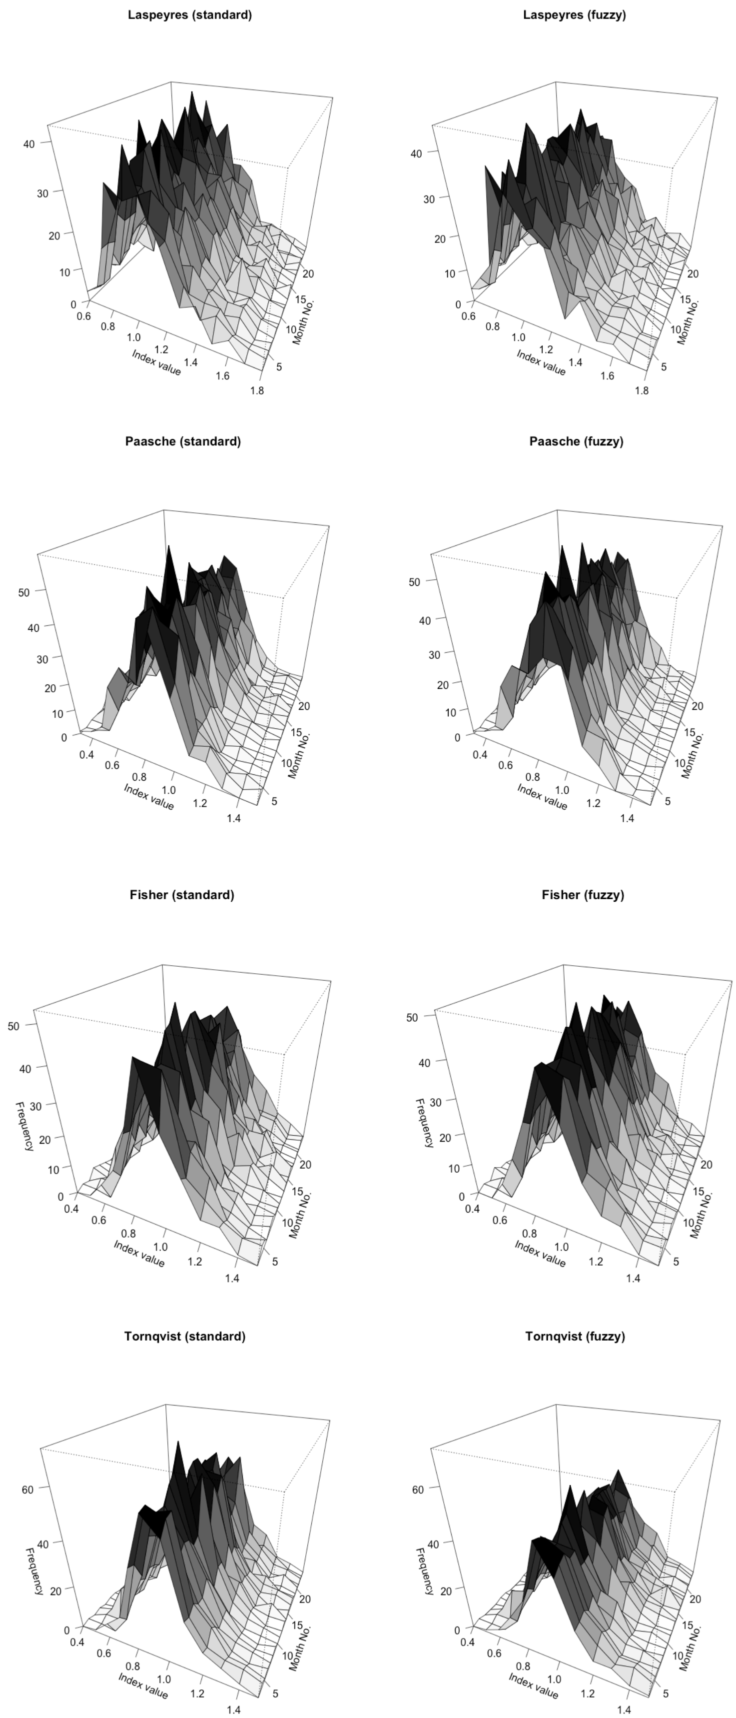

Figure 3.29: 3D surface plots (Non-chained)

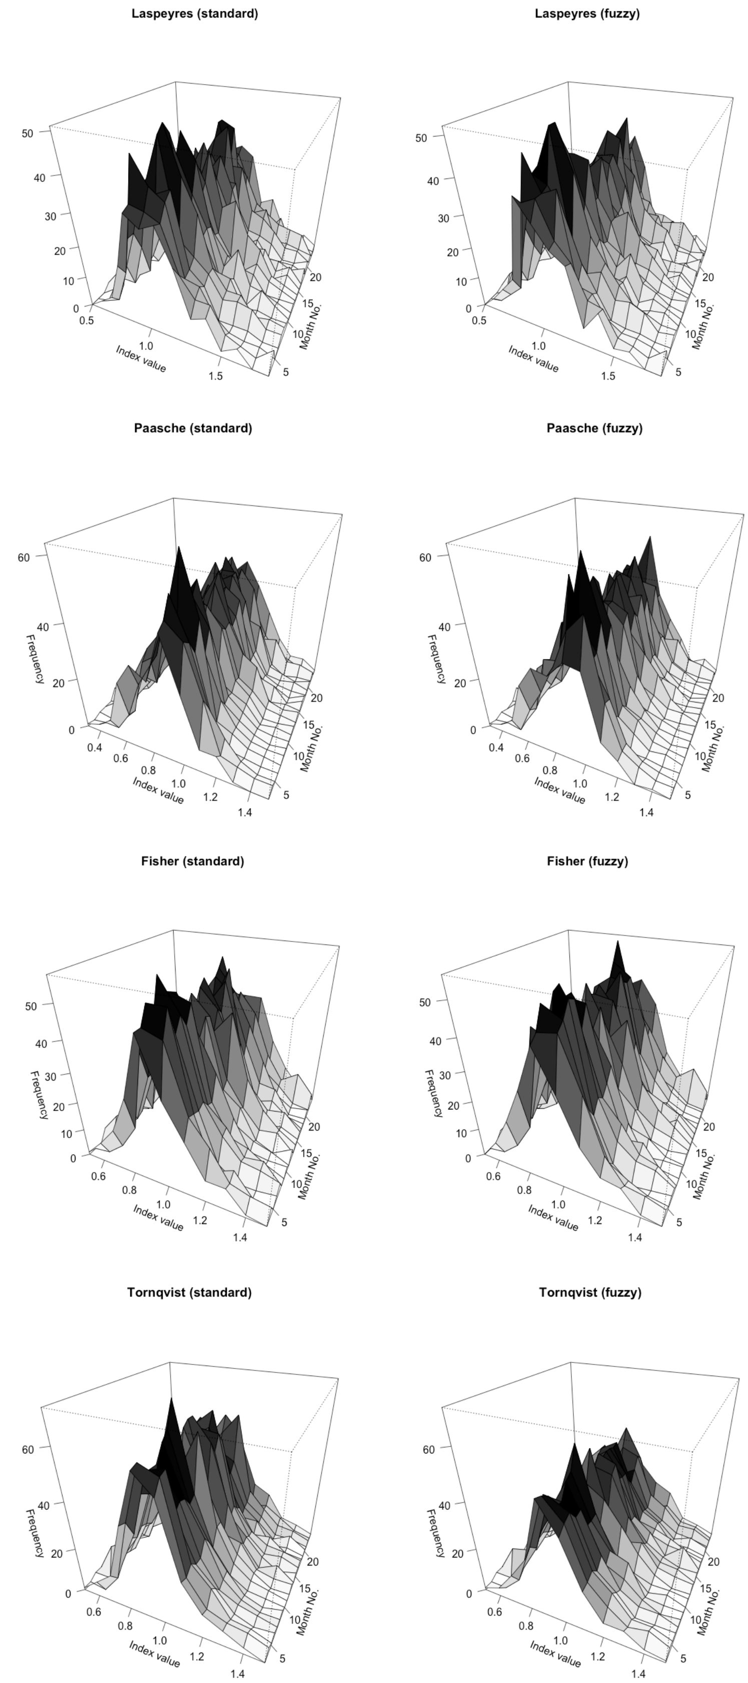

Figure 3.30: 3D surface plots (Chained)

Conclusion

This thesis used nonparametric revealed preference methods to derive new tests for consistency with models of consumer behaviour, and discuss the implications for welfare analysis.

Chapter 1 demonstrated how to conduct revealed preference analysis when prices, and hence budget constraints, are only partially observed due to rationing or corner solutions. This chapter extended the revealed preference results of Crawford and Polisson (2015), derived for the static case, to dynamic settings, in which consumers made consumption plans (Browning (1989) life cycle model) and stored goods over time (Crawford (2010), Demuynck and Verriest (2013)). Necessary and sufficient conditions for consistency with intertemporal models are derived, which do not require the researcher to distinguish between corner solutions and unavailability of the good, or to impute prices. These revealed preference conditions are also exhaustive, meaning that knowing which reason(s) resulted in the missing prices would not impose additional restrictions on the data. Simulated examples demonstrated that even if actual prices are partially observed, the revealed preference conditions are still falsifiable due to the conditions imposed on the support prices and discounted shadow prices.

Chapter 2 investigated the validity of using reported happiness measures as proxies of utility or social welfare, by testing for consistency between revealed and reported preference orderings in Japanese household survey data. To test for consistency with revealed preference, the preference orderings given by happiness numbers were included into standard revealed preference tests as additional constraints. Household behaviour that was utility-maximising and also satisfied the additional constraints would be consistent with their happiness measures. Although the expenditure behaviour of most households is consistent with standard models of utility maximisation, it is generally inconsistent with the preference ordering given by their reported happiness. This inconsistency is likely due to reporting error in the happiness measure, and suggests that taking happiness measures as an indicator or proxy for revealed preferences or utility is a very strong assumption in economic contexts.

Chapter 3 explored the effect of price inattention on inflation misperceptions and cost-of-living indices, by developing a behavioural model in which consumers only notice price changes above a certain threshold. This model is characterised by short-term memory, an asymmetric response to overspending and underspending, and a price cutoff, which can be good-specific or common to all goods. These features capture the notions of anchoring, partial adjustment towards the truth, and framing effects, all of which are part of behavioural decision-making. Necessary and sufficient conditions for data to be consistent with this price-inattention model are derived, which are easy to test given reasonable restrictions on the parameters. A data application, using supermarket scanner data, demonstrates that this model generates plausible results; in particular, consumers have more accurate perceptions of inflation during periods of high or volatile inflation, but may substantially misperceive inflation when it is low. These results have important implications for conducting welfare analysis when consumers are not fully attentive to price changes.

References

Afriat, S. N. (1967). The construction of utility functions from expenditure data. International Economic Review, 8(1), 67-77.

Afriat, S.N. (1972). Efficiency estimates of production functions. International Economic Review, 13, 568—598.

Aguiar, V. (2015). Stochastic choice and fuzzy attention. Working paper.

Antonides, G., Heijman, W., and Schouten, M. (2006). Perceived inflation and actual price changes in The Netherlands. In Heterodox views on economics and the economy of the global society, Volume 1. Wageningen Academic Publishers.

Bank of England (2010). Public attitudes to inflation and monetary policy.

Basu, K. (1997). Why are so many goods priced to end in nine? And why this practice hurts the producers. Economics Letters, 54(1), 41-44.

Becker, G. S., and Murphy, K. M. (1988). A theory of rational addiction. The Journal of Political Economy, 675-700.

Bernal, N., van Soest, A., and Vermeulen, F. (2014). Is utility informative about happiness? Discussion chapter, available at http://udep.edu.pe/cceeee/files/2014/07/3D1_Bernal_al.pdf (accessed 20 October 2016).

Bewley, T. (1977). The permanent income hypothesis: A theoretical formulation. Journal of Economic Theory, 16(2), 252-292.

Biau, O., Dieden, H., Ferrucci, G., Friz, R., and Lindén, S. (2010). Consumers' quantitative inflation perceptions and expectations in the euro area: An evaluation. Mimeo.

Brachinger, H. W. (2006). Euro or "teuro"?: The euro-induced perceived inflation in Germany. Université de Fribourg.

Bronars, S.G. (1987). The power of nonparametric tests of preference maximization. Econometrica, 55(3), 693—698.

Browning, M. (1989). A nonparametric test of the life cycle rational expectations hypothesis. International Economic Review, 979-992.

Browning, M. and Chiappori, P. A. (1998). Efficient intra-household allocations: A general characterization and empirical tests. Econometrica, 66, 1241–1278.

Cherchye, L., De Rock, B., and Vermeulen, F. (2007). The collective model of household consumption: A nonparametric characterization. Econometrica, 75, 553-574.

Cherchye, L., De Rock, B., and Vermeulen, F. (2008). The revealed preference approach to collective consumption behavior: Testing, recovery and welfare analysis. Center Discussion chapter No. 2007–73, University of Tilburg.

Cherchye, L., De Rock, B., and Vermeulen, F. (2009). Opening the black box of intra-household decision-making: Theory and non-parametric empirical tests of general collective consumption models. Journal of Political Economy, 117, 1074–1104.

Cherchye, L., Demuynck, T., De Rock, B. (2011a). Testable implications of general equilibrium models: an integer programming approach. Journal of Mathematical Economics, 47, 564-575.

Cherchye, L., De Rock, B., and Vermeulen, F. (2011b). The revealed reference approach to collective consumption behavior: Testing and sharing rule recovery. Review of Economic Studies, 78, 176–198.

Cherchye, L., Demuynck, T., and Rock, B. D. (2011c). Revealed preference analysis of non-cooperative household consumption. The Economic Journal, 121(555), 1073-1096.

Cherchye, L., De Rock, B., and Platino, V. (2013). Private versus public consumption within groups: Testing the nature of goods from aggregate data. Economic Theory, 54(3), 485-500.

Cherchye, L., Demuynck, T., and De Rock, B. (2015). Is utility transferable? A revealed preference analysis. Theoretical Economics, 10(1), 51-65.

Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: theory and evidence. American Economic Review 99(4), 1145–1177.

Crawford, I. (2010). Habits revealed. The Review of Economic Studies, 77(4), 1382-1402.

Crawford, I., and Polisson, M. (2015). Demand analysis with partially observed prices (No. W15/16). IFS Working Papers.

Deaton, A. (1992). Understanding Consumption. Oxford University Press.

Deaton, A., and Irish, M. (1984). Statistical models for zero expenditures in household budgets. Journal of Public Economics, 23(1-2), 59-80.

Deaton, A. (1988). Quality, quantity, and spatial variation of price. The American Economic Review, 418-430.

Demuynck, T., and Verriest, E. (2013). I'll never forget my first cigarette: A revealed preference analysis of the 'Habits as Durables' model. International Economic Review, 54(2), 717-738.

Di Tella, R., MacCulloch, R. J., and Oswald, A. J. (2001). Preferences over inflation and unemployment: Evidence from surveys of happiness. The American Economic Review, 91(1), 335-341.

Di Tella, R., MacCulloch, R. J., and Oswald, A. J. (2003). The macroeconomics of happiness. Review of Economics and Statistics, 85(4), 809-827.

Ferrer-i-Carbonell, A., and Van Praag, B. (2002). The subjective costs of health losses due to chronic diseases. An alternative model for monetary appraisal. Health Economics, 11(8), 709-722.

Ferrer i Carbonell, A., and Frijters, P. (2004). How important is methodology for the estimates of the determinants of happiness? The Economic Journal, 114(497), 641-659.

Ferrer-i-Carbonell, A. (2013). Happiness economics. SERIEs, 4(1), 35-60.

Finkelstein, A., Luttmer, E. F., and Notowidigdo, M. J. (2013). What good is wealth without health? The effect of health on the marginal utility of consumption. Journal of the European Economic Association, 11(S1), 221-258.

Frank, R. H. (1997). The frame of reference as a public good. Economic Journal, 107(445), 1832-1847.

Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time discounting and time preference: A critical review. Journal of Economic Literature, 40(2), 351-401.

Frey, B. S., and Stutzer, A. (2002). What can economists learn from happiness research? Journal of Economic Literature, 40(2), 402-435.

Frijters, P., Johnston, D. W., and Shields, M. A. (2011). Life satisfaction dynamics with quarterly life event data. The Scandinavian Journal of Economics, 113(1), 190-211.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. The Quarterly Journal of Economics, 129(4), 1661-1710.

Groot, W., van den Brink, H. M., and Plug, E. (2004). Money for health: The equivalent variation of cardiovascular diseases. Health Economics, 13(9), 859-872.

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. Journal of Political Economy, 86, 971-987.

Headey, B. and Wearing, A. (1989). Personality, life events, and subjective well-being: Toward a dynamic equilibrium model. Journal of Personality and Social Psychology, 57, 731-739.

Kahneman, D. and Krueger, A. B. (2006). Developments in the measurement of subjective well-being. The Journal of Economic Perspectives, 20(1), 3-24.

Keen, M. (1986). Zero expenditures and the estimation of Engel curves. Journal of Applied Econometrics, 1(3), 277-286.

Konüs, A. A. (1924). The problem of the true index of the cost of living. Econometrica, 7, 0–29.

Layard, R., Mayraz, G., and Nickell, S. (2008). The marginal utility of income. Journal of Public Economics, 92(8), 1846-1857.

Levinson, A. (2009). Valuing air quality using happiness data. Mimeo, Georgetown University.

Luechinger, S. (2009). Valuing air quality using the life satisfaction approach. The Economic Journal, 119(536), 482-515.

Meghir, C., and Robin, J. M. (1992). Frequency of purchase and the estimation of demand systems. Journal of Econometrics, 53(1), 53-85.

Neary, J. P., and Roberts, K. W. (1980). The theory of household behaviour under rationing. European Economic Review, 13(1), 25-42.

O'Donoghue, J. (2007). Inflation–experience and perceptions. Economic and Labour Market Review, 1(1), 33-39.

Oswald, A. J. (1997). Happiness and economic performance. The Economic Journal, 107(445), 1815-1831.

Oswald, A. J., and Powdthavee, N. (2008). Death, happiness, and the calculation of compensatory damages. The Journal of Legal Studies, 37(S2), S217-S251.

Plug, E. J., and Van Praag, B. M. (1995). Family equivalence scales within a narrow and broad welfare context. Journal of Income Distribution, 4(2), 1-1.

Polisson, M., and Quah, J. K. H. (2013). Revealed preference in a discrete consumption space. American Economic Journal: Microeconomics, 5(1), 28-34.

Ryder, H. E., and Heal, G. M. (1973). Optimal growth with intertemporally dependent preferences. The Review of Economic Studies, 40(1), 1-31.

Schepelmann, P., Goossens, Y., and Makipaa, A. (2010). Towards sustainable development: Alternatives to GDP for measuring progress. Wuppertal Spezial 42. Wuppertal Institute for Climate, Environment, and Energy, Wuppertal, Germany. Available at http://wupperinst.org/en/publications/details/wi/a/s/ad/1113/ (accessed 18 October 2016).

Selten, R. (1991). Properties of a measure of predictive success. Mathematical Social Sciences, 21, 153-167.

Spinnewyn, F. (1981). Rational habit formation. European Economic Review, 15(1), 91-109.

Stiving, M. (2000). Price-endings when prices signal quality. Management Science, 46(12), 1617-1629.

Stutzer, Alois and Frey, Bruno S. (2012). Recent Developments in the Economics of Happiness: A Selective Overview. IZA Discussion chapter No. 7078. Available at SSRN: https://ssrn.com/abstract=2192854.

Tversky, A., and Kahneman, D. (1975). Judgment under uncertainty: Heuristics and biases. In Utility, probability, and human decision making (pp. 141-162). Springer Netherlands.

Tversky, A., and Kahneman, D. (1985). The framing of decisions and the psychology of choice. In Environmental Impact Assessment, Technology Assessment, and Risk Analysis (pp. 107-129). Springer Berlin Heidelberg.

Van den Berg, B., and Ferrer-i-Carbonell, A. (2007). Monetary valuation of informal care: The well-being valuation method. Health economics, 16(11), 1227-1244.

Van Praag, B. M., Frijters, P., and Ferrer-i-Carbonell, A. (2003). The anatomy of subjective well-being. Journal of Economic Behavior and Organization, 51(1), 29-49.

Varian, H. R. (1982). The nonparametric approach to demand analysis. Econometrica, 945-973.

Varian, H. (1985). Non-parametric analysis of optimizing behavior with measurement error. Journal of Econometrics, 30, 445-458.

Varian, H. R. (1988). Revealed preference with a subset of goods. Journal of Economic Theory, 46(1), 179-185.

Varian, H. (1990). Goodness-of-fit in optimizing models. Journal of Econometrics, 46, 125—140.

Zadeh, L. A. (1965). Fuzzy sets. Information and control, 8(3), 338-353.