

When Does Static Willingness to Pay Mislead? A Framework for Dynamic Hedonic Valuation*

Josephine Auer[†]

MIT

May 19, 2026

Abstract

Many policy counterfactuals depend on how consumers value product attributes such as sugar, sodium, caffeine, alcohol, or emissions. Standard hedonic and differentiated-products models interpret these valuations statically. That interpretation is restrictive when attributes are habit forming: observed prices then reflect both contemporaneous marginal value and the continuation value generated by current consumption. I develop a nonparametric revealed-preference framework for dynamic hedonic valuation, deriving necessary and sufficient conditions for rationalising observed prices and choices. Using household scanner data on cereal purchases, I show that the hedonic representation places real restrictions on prices, while habit formation improves behavioural coherence conditional on that representation. The results provide a diagnostic for when static attribute valuation is justified and when prices reveal more than contemporaneous marginal values.

Keywords: Hedonic valuation, habit formation, identification, revealed preference, intertemporal choice, nonparametric analysis

JEL Classification: D11, D12, D90, C61, C14, L66

*I am especially grateful to Nikhil Agarwal, Isaiah Andrews, Miguel Ballester, Abhijit Banerjee, Nano Barahona, Colin Camerer, Ian Crawford, Glenn Ellison, Drew Fudenberg, Eric Gao, Matt Polisson, Augustus Smith, and Chen Zhen for helpful comments and discussions.

[†]E-mail: joauer@mit.edu

1 Introduction

Taxes, regulations, and product standards often target characteristics rather than products themselves. Soda taxes target sugar, nutrition policy targets sodium and other dietary attributes, tobacco and alcohol policy target addictive substances, and environmental standards target emissions or energy efficiency. Applied economists routinely value such characteristics using hedonic or differentiated-products demand models, and the resulting willingness-to-pay (WTP) measures are then used to interpret welfare and policy counterfactuals.¹ The usual interpretation is static: a price or demand coefficient attached to a characteristic is read as a contemporaneous marginal valuation.

This static interpretation is restrictive when the characteristic itself is habit forming. If consuming an additional gram of sugar today lowers tomorrow’s marginal utility through fatigue, satiation, or health-related costs, then the observed willingness to pay for that gram combines a current benefit with a negative continuation value. In that case, a static researcher understates contemporaneous WTP by treating the whole price as a current marginal value. If current consumption instead raises future marginal utility through reinforcement or craving, the static interpretation moves in the opposite direction. Static WTP objects may therefore confound current tastes, continuation values, and the future preference states relevant for welfare counterfactuals.

The central question of this paper is when characteristic-based valuation is coherent, and when it requires a dynamic rather than static interpretation. Existing models of habit formation are typically formulated over *goods*, not characteristics. A large literature studies intertemporal dependence in the consumption of cigarettes, alcohol, food, and digital products by modelling utility as a function of past and current quantities (Becker and Murphy, 1988; Gruber and Kőszegi, 2004; Demuyne and Verriest, 2013; Crawford, 2010; Allcott et al., 2022). These models characterise dynamic dependence in product quantities, but they do not determine whether prices admit a coherent interpretation as valuations of the underlying attributes that policy often targets. When valuation is conducted in characteristics space, the dynamic structure must also be specified in characteristics space.

I provide a nonparametric revealed-preference (RP) framework for evaluating the static interpretation of characteristic-based valuation. The framework asks whether observed prices and choices can be rationalised by preferences over current characteristics and lagged consumption of a habit-forming subset of those characteristics. The exercise is consumer by consumer, so valuations over characteristics, discounting, and habit formation may differ freely across households. The goal is not to estimate a parametric demand system, but to derive the exact restrictions that dynamic characteristic-based valuation places on observables.

The first contribution is a sharp Afriat-type characterisation of dynamic hedonic rationalisability. The test asks whether there exist characteristic shadow prices and a discount factor such that two requirements hold. First, observed goods prices must admit the shadow-price representation implied by the maintained goods-to-characteristics technology. Second, observed choices must satisfy dynamic RP inequalities at those shadow prices. This extends static characteristics-based rationalisability (e.g., Blow et al. (2008)) to allow habits over characteristics, while preserving the nonparametric logic of Samuelson (1948), Houthakker (1950), Afriat (1967), and Browning (1989).

The second contribution is to show how dynamics alter the interpretation of characteristic prices. In a static hedonic model, the price of a good purchased in positive quantity equals the value of the characteristics

¹Hedonic and characteristic-based demand models are central in empirical industrial organisation and applied micro; see, e.g., Smith and Desvousges (1986), Heckman and Scheinkman (1987), Berry et al. (1995), Nevo (2001), Gibbons and Machin (2003), Bajari and Kahn (2005), and Greenstone and Gallagher (2008). For examples valuing potentially habit-forming food attributes under static hedonic preferences, see Dubois et al. (2020), Haecck et al. (2022), and Le Fur and Outreville (2022).

bundled in that good. In the dynamic model, the same price equals contemporaneous marginal valuation plus the discounted continuation value generated by habit-forming attributes. Static WTP is therefore correct only under an exclusion restriction: current consumption of the characteristic must have no future utility effect. The same issue arises in differentiated-products demand, where characteristic coefficients are often converted into WTP for policy-relevant attributes (Berry et al., 1995; Nevo, 2001). Even if such a coefficient is correctly identified in a static demand system, it is a contemporaneous welfare object only under the same exclusion restriction. Otherwise, static counterfactuals built from the coefficient can mismeasure welfare by ignoring how current consumption changes future marginal valuations.

The third contribution is to decompose empirical failure into two distinct margins. A characteristics-based valuation exercise can fail because observed prices cannot be represented through the maintained characteristics technology; this is a structural failure of the hedonic price system. It can also fail because, conditional on such shadow prices existing, observed choices are not dynamically rationalisable; this is a behavioural failure of intertemporal choice. For applied work, this distinction matters because a failed valuation exercise has different implications depending on whether the problem is the characteristics technology itself or the static interpretation of behaviour at admissible shadow prices. For example, if soda prices do not admit a shadow price for sugar under the maintained characteristics technology, then a soda-tax calculation cannot be grounded in hedonic WTP for sugar. If such shadow prices exist, but only the dynamic model rationalises behaviour, the admissible WTP object combines current utility from sugar with the future utility effects of current sugar consumption.

I complement the binary test with diagnostics that measure how severe violations are on each margin. On the structural margin, I compute the distance from observed prices to the hedonic price manifold implied by the characteristics technology. On the behavioural margin, I adapt a Critical Cost Efficiency Index (CCEI) (Afriat, 1973; Varian, 1990) to measure the smallest relaxation of the dynamic RP inequalities needed for rationalisability. These diagnostics distinguish small departures from consistency from large ones and show whether empirical discipline comes primarily from the price representation or from intertemporal choice.

I apply the framework to household-level scanner data on cold-cereal purchases. Cereal is a natural setting for the exercise because products bundle measurable nutritional and descriptive characteristics, purchases are repeated over time, and certain attributes—such as sugar and sodium—are plausibly habit-forming and relevant for nutrition policy. More generally, the same framework applies to any setting with repeated purchases, observed expenditures, and measured product characteristics.

The application yields three findings. First, moving from goods to characteristics space sharply lowers raw rationalisability rates because characteristic-based models impose real restrictions on the price system. Second, the associated distance measures show that many of these structural violations are modest in dollar terms, so binary rejection alone overstates their economic severity. Third, conditional on the hedonic price representation being admissible, allowing for habits improves behavioural coherence relative to static characteristics models: it raises raw rationalisability rates for some households and moves the distribution of CCEI values closer to one more generally. These patterns show that the framework is not simply a more flexible fit criterion: it diagnoses when static characteristic-based valuation is empirically defensible and when a dynamic interpretation is additionally required.

The results clarify when prices and choices support characteristic-based valuation, when that valuation requires a dynamic interpretation, and what this implies for policy counterfactuals built from static WTP. If observed prices do not admit a characteristics-based representation, no static or dynamic hedonic interpretation is coherent under the maintained technology. If the price representation is admissible but static

separability fails, then characteristic prices can no longer be read purely as contemporaneous WTP. The continuation component is generally set identified rather than point identified, so the framework is best understood as a diagnostic of admissibility and interpretation rather than as a full structural estimator of welfare primitives.

The remainder of the paper is organised as follows. Section 2 introduces the dynamic hedonic model, derives the valuation wedge, and states the RP characterisation. Section 3 situates the result within the related RP literature by showing that static hedonic and dynamic goods-based models arise as special cases. Section 4 applies the framework to household scanner data on cereal purchases.

Related work

This paper first relates to the hedonic valuation and differentiated-products demand literatures following Rosen (1974) and Berry et al. (1995) (henceforth BLP).² Empirical work in this tradition often uses observed prices or demand coefficients to value product characteristics that enter welfare and policy counterfactuals. My point of departure is the static interpretation of those values. When some attributes are habit forming, observed prices or demand coefficients can confound current tastes with continuation values, so the economic interpretation of WTP for characteristics becomes an empirical question rather than a maintained one. This distinction is first-order when WTP estimates are used to evaluate taxes, subsidies, or standards that target the characteristics themselves.

The interpretation problem studied here is distinct from the familiar price-endogeneity problem in differentiated-products demand. In BLP-style models, instruments and supply-side restrictions are used to separate demand from marginal costs, markups, and unobserved product quality. My exercise instead conditions on the prices faced by a household and asks whether those prices and choices admit a characteristic-based welfare interpretation. The two issues are complementary: even if a demand system has recovered a static WTP coefficient for sugar, sodium, or emissions, that coefficient need not be a primitive taste parameter if current consumption of the characteristic changes future marginal utility. In that sense, the paper adds an interpretation diagnostic to structural hedonic and demand approaches such as Bajari and Benkard (2005): it asks whether recovered WTP objects can be interpreted statically, or whether habit-forming characteristics require a dynamic interpretation.

The paper is also related to the rational addiction literature initiated by Becker and Murphy (1988).³ That literature shows why current consumption may affect future marginal utility, but it is formulated at the level of goods. The distinction matters when persistence is more naturally attached to attributes such as sugar, sodium, nicotine, or caffeine than to the composite products that bundle them (cf. Koob and Moal, 1997). Relative to empirical work on addiction over nutrient profiles (e.g. Richards and Patterson, 2006), my approach replaces functional-form restrictions with a nonparametric RP characterisation and asks when habits over characteristics are consistent with observed prices and choices.

Methodologically, the paper contributes to two nonparametric RP literatures that have developed separately: tests of preferences over characteristics (Blow et al., 2008) and RP analyses of intertemporal dependence (Crawford, 2010), both rooted in the foundational work of Afriat (1967), Diewert (1973), and Varian (1982). I build on these tools to answer a valuation question that neither literature addresses on its own: when do observed prices and choices support a dynamic interpretation of WTP for characteristics?

²Rosen's (1974) hedonic method has been widely applied across housing, environmental economics, health, education, and industrial organisation; see, for example, Smith and Desvousges (1986), Gibbons and Machin (2003), Bajari and Kahn (2005), and Greenstone and Gallagher (2008).

³Empirical applications have largely focused on goods such as alcohol, cigarettes, caffeine, and illicit drugs; see, for example, Becker et al. (1994), Grossman et al. (1998), Gruber and Kőszegi (2004), and Demuyne and Verriest (2013).

Answering this question requires both a goods-to-characteristics price representation and intertemporal rationalisability at the implied shadow prices. Static hedonic RP tests and dynamic goods-based RP tests therefore arise as limiting cases, but the interaction between the two restrictions is the object of interest. This interaction yields a novel taxonomy of empirical failure and tractable diagnostics for its severity.

Finally, the framework is distinct from models of inventory behaviour and stockpiling in consumer demand (Hendel and Nevo, 2006). Storage can generate persistence in purchases under time-separable preferences, whereas habit formation in my setting operates through state dependence in the flow utility function over characteristics. The empirical framework therefore treats persistence induced by habits as conceptually different from persistence induced by intertemporal substitution, even though distinguishing the two in purchase data requires a maintained link between purchases and the underlying consumption state.

Taken together, the paper shows when characteristic-based valuation is admissible, and when the resulting valuations require a dynamic rather than static interpretation. Moving to characteristics space imposes demanding structural restrictions on the admissible price system; allowing for dynamics changes the set of admissible behavioural patterns and the interpretation of shadow prices conditional on satisfying those restrictions.

2 Model

I begin by formalising a dynamic hedonic environment in which goods map into characteristics and habits attach to a subset of those characteristics. I observe a single consumer for $t = 1, \dots, T$, with purchases $\mathbf{x}_t \in \mathbb{R}_+^K$ and present-value prices $\boldsymbol{\rho}_t \in \mathbb{R}_+^K$. Observed goods prices are taken as given from the consumer’s perspective, so the object of the analysis is not price determination but whether observed choices can be rationalised. Goods map into J measured characteristics via a time-invariant linear technology $\mathbf{z}_t = \mathbf{A}\mathbf{x}_t$ following Gorman (1956), where \mathbf{A} is a $J \times K$ matrix (typically $J < K$).⁴ I treat \mathbf{A} as known and stable over time; it captures objectively defined product attributes, while valuation is encoded in preferences. Although the model is written for one consumer, the empirical RP exercise is applied household by household, so valuations over characteristics, discounting, and habit formation are all allowed to vary freely across consumers.

I partition characteristics as $\mathbf{z}_t = ((\mathbf{z}_t^c)', (\mathbf{z}_t^a)')'$, where $\mathbf{z}_t^c \in \mathbb{R}^{J_1}$ are non-habit-forming characteristics and $\mathbf{z}_t^a \in \mathbb{R}^{J_2}$ are habit-forming characteristics, with $J_1 + J_2 = J$. The analyst defines this partition. Crucially, this formulation allows intertemporal dependence to operate at the level of attributes rather than goods, so persistence in behaviour need not be attributed to non-habit-forming components bundled within a product. Although preferences are defined over characteristics, choice and budget constraints remain in goods space.

To clarify the economic content of the model, it is useful to fix ideas with a simple example. Suppose the consumer chooses between two cereal products, where each good bundles two measurable attributes: a contemporaneous “taste” characteristic (e.g., nuttiness) and a habit-forming “sensory” characteristic (e.g., salt or sugar intensity). The key modelling choice is that habits attach to the latter attribute rather than to the cereal good itself: consuming a high-intensity product today can change tomorrow’s marginal value of intensity (due to sensory fatigue or craving), even if the consumer switches cereal product. In this example, the model’s structural content is that observed goods prices must be representable as shadow values on

⁴I adopt a linear transformation from goods to characteristics space as it is the most widely used specification. Most results extend to a non-linear setting where $\mathbf{z} = \mathbf{F}(\mathbf{x})$ is increasing and strictly concave; see Appendix B for an analogue of the consistency definition. The key difference to the linear case arises in the marginal product: whereas $\partial \mathbf{z} / \partial \mathbf{x} = \mathbf{A}'$ is constant in the linear model, the marginal product varies with demand under a non-linear transformation.

attributes given $\mathbf{z}_t = \mathbf{A}\mathbf{x}_t$, while its behavioural content is that those shadow values must admit a concave, dynamically consistent utility representation.

Preferences are represented by a felicity (i.e., flow utility) function $u : \mathbb{R}^{J+J_2} \rightarrow \mathbb{R}$ that depends on current characteristics and one lag of the habit-forming subset, $u(\mathbf{z}_t^c, \mathbf{z}_t^a, \mathbf{z}_{t-1}^a)$, as in the one-lag “short memory habits” specification of Boyer (1978, 1983) and Becker et al. (1994). The multi-lag extension is straightforward and deferred to Online Appendix F. I assume quasi-linearity in an outside good y_t with unit price, as is standard in empirical IO and hedonic demand models for narrow product categories (Berry et al., 1995; Nevo, 2001). I also maintain local non-satiation, concavity, and superdifferentiability of u ; I impose no further sign or monotonicity restrictions on the habit-forming components. I refer to this environment as the *habits-over-characteristics* model.

The consumer chooses $\{(\mathbf{x}_t, y_t)\}_{t=1}^T$ to solve

$$\max_{\{(\mathbf{x}_t, y_t)\}_{t=1}^T} \sum_{t=1}^T \beta^{t-1} (u(\tilde{\mathbf{z}}_t) + y_t) \quad \text{s.t.} \quad \sum_{t=1}^T \boldsymbol{\rho}'_t \mathbf{x}_t + \sum_{t=1}^T \beta^{t-1} y_t = W, \quad \tilde{\mathbf{z}}_t = \tilde{\mathbf{A}} \tilde{\mathbf{x}}_t, \quad (1)$$

where $\beta \in (0, 1]$ is a discount factor and W is present-value lifetime wealth. I use the augmented notation

$$\tilde{\mathbf{z}}_t := ((\mathbf{z}_t^c)', (\mathbf{z}_t^a)', (\mathbf{z}_{t-1}^a)')', \quad \tilde{\mathbf{x}}_t := (\mathbf{x}'_t, \mathbf{x}'_{t-1})', \quad \tilde{\mathbf{A}} := \begin{pmatrix} \mathbf{A} & \mathbf{0}_{J \times K} \\ \mathbf{0}_{J_2 \times K} & \mathbf{A}^a \end{pmatrix}, \quad (2)$$

so $\tilde{\mathbf{z}}_t \in \mathbb{R}^{J+J_2}$, $\tilde{\mathbf{x}}_t \in \mathbb{R}^{2K}$, and $\tilde{\mathbf{A}}$ is a $(J + J_2) \times 2K$ block matrix. Throughout, I treat \mathbf{x}_0 (equivalently, the initial habit stock \mathbf{z}_0^a) as fixed and exogenous. Because the model has one lag and a finite horizon, there is no continuation term beyond T ; equivalently, I set $\boldsymbol{\pi}_{T+1}^1 \equiv \mathbf{0}$ by convention. The key question is whether the observables $\{(\boldsymbol{\rho}_t, \mathbf{x}_t)\}_{t=1}^T$ can be rationalised by (1), and, if so, what testable restrictions this imposes on prices and choices.

2.1 Consistency

I now ask whether the observed data can be rationalised by optimising behaviour under the habits-over-characteristics model. This notion of consistency encompasses both the existence of a hedonic shadow-price representation (a structural requirement) and the coherence of intertemporal behaviour conditional on that representation. Here “structural” does not refer to price determination. Rather, it refers to whether the maintained goods-to-characteristics technology can rationalise observed prices through a corresponding system of characteristic shadow prices.

Definition 2.1 (Consistency). The data $\{(\boldsymbol{\rho}_t, \mathbf{x}_t)\}_{t=1}^T$ are *consistent* with the one-lag habits-over-characteristics model for given \mathbf{A} if there exist $\beta \in (0, 1]$ and a locally non-satiated, concave, superdifferentiable felicity function u such that $\{\mathbf{x}_t\}_{t=1}^T$ solves (1).

The following lemma provides necessary and sufficient conditions for consistency.

Lemma 2.1 (Consistency Conditions). The data $\{\boldsymbol{\rho}_t, \mathbf{x}_t\}_{t=1}^T$ are consistent with the one-lag habits-over-characteristics model for a given technology matrix \mathbf{A} if and only if there exist a locally non-satiated, concave, superdifferentiable utility function $u(\cdot)$ and a discount factor $\beta \in (0, 1]$ such that for

all $t \in \{1, \dots, T\}$,

$$\rho_t \geq \mathbf{A}'\boldsymbol{\pi}_t^0 + (\mathbf{A}^a)'\boldsymbol{\pi}_{t+1}^1, \quad (\star)$$

with equality for all goods k such that $x_t^k > 0$, where the discounted shadow prices are

$$\boldsymbol{\pi}_t^0 = \beta^{t-1} \begin{bmatrix} \partial_{\mathbf{z}_t^c} u(\tilde{\mathbf{z}}_t) \\ \partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_t) \end{bmatrix}, \quad (3)$$

$$\boldsymbol{\pi}_t^1 = \beta^{t-1} \partial_{\mathbf{z}_{t-1}^a} u(\tilde{\mathbf{z}}_t), \quad (4)$$

and $\tilde{\mathbf{z}}_t = \tilde{\mathbf{A}}\tilde{\mathbf{x}}_t$, with $\boldsymbol{\pi}_{T+1}^1 \equiv \mathbf{0}$.

Proof: See Appendix A. \square

2.2 Dynamic wedges in willingness to pay

Lemma 2.1 shows why habits can matter for characteristic-based valuation. The shadow price $\boldsymbol{\pi}_t^0$ is the discounted marginal valuation of contemporaneous characteristics, while $\boldsymbol{\pi}_t^1$ captures the marginal utility effect of past consumption of habit-forming characteristics. For any good purchased in positive quantity, condition (\star) implies that its price equals the sum of current characteristic values and the discounted continuation value generated by habit-forming attributes.

This decomposition creates a wedge between static and dynamic WTP. In a static hedonic model, the price of a good is interpreted as the value of the characteristics bundled in that good. With habits, the same price also internalises the future utility effect of current consumption. For example, when current consumption lowers future marginal utility (i.e., $\partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_{t+1}) < 0$), goods prices satisfy

$$\rho_t^k = \mathbf{a}'_k \boldsymbol{\pi}_t^0 + \mathbf{a}'_k \boldsymbol{\pi}_{t+1}^1 < \mathbf{a}'_k \boldsymbol{\pi}_t^0,$$

so ignoring intertemporal dependence understates contemporaneous WTP for goods with negatively reinforcing attributes. If current consumption instead raises future marginal utility through reinforcement or craving, the inequality reverses. The static interpretation is therefore correct only when the continuation component is zero.

This wedge matters when WTP estimates are used to interpret attribute-targeted policy. Consider a tax that raises product prices in proportion to sugar or sodium content. If the narrow question is whether a higher price lowers current purchases, the decomposition of WTP into current and continuation components may not be central. The decomposition matters once the WTP estimate is used as a welfare object. A static interpretation treats the observed characteristic price as current marginal surplus from the taxed characteristic. In the dynamic model, the same object is a lifetime marginal value: it includes current utility and the continuation value created by today's exposure. A policy analyst may therefore predict a demand response while mischaracterising the welfare value of that response.

The same logic applies to differentiated-products demand models (Berry et al., 1995; Nevo, 2001). In a static model with product characteristics, a coefficient on sugar, sodium, emissions, or energy efficiency is often converted into WTP and used for counterfactual welfare analysis. If preferences are dynamically dependent over those characteristics, that coefficient is a static welfare object only under an exclusion restriction: current consumption of the characteristic must not affect future marginal utility. If the restriction

fails, standard instruments and supply-side restrictions may still recover a static demand coefficient, but that coefficient is not the contemporaneous WTP object required for static welfare counterfactuals.

This point is distinct from the usual equilibrium price problem in BLP-style demand estimation. There, instruments and supply-side restrictions are used to separate demand from marginal costs, markups, and unobserved product quality. Here, prices are taken as part of the household’s choice environment, and the question is whether those prices and choices support a characteristic-based welfare interpretation, and whether that interpretation is static or dynamic.

The implication extends to BLP-style supply-side counterfactuals. In a static Bertrand model, firms’ first-order conditions use current demand elasticities to map prices into markups and marginal costs (Berry et al., 1995). But if consumers are dynamic in characteristics, the relevant demand response reflects not only the current utility effect of a price change but also how price-induced changes in current consumption affect future states and continuation values. This paper does not estimate such a dynamic demand system. Rather, it provides a diagnostic for when static WTP and static demand-side inputs may be missing the dynamic margin, even when firms themselves are modelled as static.

For an applied demand researcher, the practical role of the test is as a diagnostic pre-test and interpretation check. Given a maintained characteristics technology, the first step is to ask whether observed prices support characteristic shadow values. Conditional on that representation, the second step is to ask whether behaviour at those shadow values is consistent with static separability, or whether allowing current characteristics to affect future marginal utility improves rationalisability. The distance-based diagnostics developed below are useful even when the model rejects: large violations warn against treating recovered characteristic WTP as a static welfare primitive, while small violations support the static model as a useful approximation. The framework is therefore complementary to structural demand estimation: it clarifies when the WTP objects recovered by such models can be read statically, and when they require a dynamic interpretation.

In sum, the model’s empirical content is governed by two restrictions: a *structural* requirement that observed goods prices admit characteristic shadow prices consistent with the maintained goods-to-characteristics technology, and a *behavioural* requirement that these shadow prices be consistent with concave, dynamically stable preferences.

2.3 Afriat conditions for habits-over-characteristics

I now state the central theoretical result of the paper. It shows that the model has two distinct sources of empirical content: a structural requirement that observed goods prices admit characteristic shadow prices given \mathbf{A} , and a behavioural requirement that these shadow prices be consistent with concave, dynamically stable preferences.

Theorem 2.1. The following statements are equivalent:

- (A) The data $\{\boldsymbol{\rho}_t; \mathbf{x}_t\}_{t \in \{1, \dots, T\}}$ are consistent with the one-lag habits model for given technology \mathbf{A} .
- (B) There exist T J -vector discounted shadow prices $\{\boldsymbol{\pi}_t^0\}_{t \in \{1, \dots, T\}}$, T J_2 -vector discounted shadow

prices $\{\boldsymbol{\pi}_t^1\}_{t \in \{1, \dots, T\}}$ and a discount factor $\beta \in (0, 1]$ such that,

$$0 \leq \sum_{m=1}^M \tilde{\boldsymbol{\pi}}'_{t_m} (\tilde{\mathbf{z}}_{t_{m+1}} - \tilde{\mathbf{z}}_{t_m}) \quad \forall M \geq 2, \forall (t_1, \dots, t_M) \in \{1, \dots, T\}^M, t_{M+1} = t_1 \quad (\text{B1})$$

$$\rho_t^k \geq \mathbf{a}'_k \boldsymbol{\pi}_t^0 + \mathbf{a}'_k \boldsymbol{\pi}_{t+1}^1 \quad \forall k, t \in \{1, \dots, T\} \quad (\text{B2})$$

$$\rho_t^k = \mathbf{a}'_k \boldsymbol{\pi}_t^0 + \mathbf{a}'_k \boldsymbol{\pi}_{t+1}^1 \quad \text{if } x_t^k > 0, \forall k, t \in \{1, \dots, T\} \quad (\text{B3})$$

where \mathbf{a}_k is the J -vector corresponding to the k -th column of \mathbf{A} , \mathbf{a}_k^a is the J_2 -vector corresponding to the last J_2 rows of the k -th column of \mathbf{A} , and $\tilde{\boldsymbol{\pi}}_t := \frac{1}{\beta^{t-1}} [\boldsymbol{\pi}_t^0, \boldsymbol{\pi}_t^1]'$, with $\boldsymbol{\pi}_{T+1}^1 \equiv \mathbf{0}$.

Proof: See Appendix A. \square

Theorem 2.1 delivers a complete RP characterisation of the one-lag habits-over-characteristics model: the data are rationalisable if and only if there exist shadow prices and a discount factor satisfying conditions (B1)–(B3). When such objects exist, one can construct a concave, locally non-satiated utility function over characteristics that rationalises observed choices; when they do not, no such representation is possible. Given the boundary convention stated above, the final-period pricing restriction is simply (B2)–(B3) with $\boldsymbol{\pi}_{T+1}^1 = \mathbf{0}$. The theorem therefore delivers a sharp, nonparametric test of dynamic consistency in characteristics space.

It is useful to interpret the economic content of the three conditions. Condition (B1) imposes cyclical monotonicity on the shadow prices. Economically, it rules out “cycles” in revealed marginal valuations over augmented characteristic bundles: there should be no sequence of observed trades in characteristics space that would allow a costless improvement by returning to the starting point. Formally, cyclical monotonicity is equivalent to concavity of the instantaneous utility function (Rockafellar, 1970). This condition is precisely the behavioural discipline of the model.

Conditions (B2) and (B3) impose the structural pricing restrictions implied by the habits-over-characteristics model. They require that observed goods prices be representable as linear combinations of contemporaneous and forward-looking shadow prices. Economically, current goods prices must internalise both current marginal utility from characteristics and the continuation value induced by habit formation. In the running cereal example, the price of a high salt or sugar cereal must reflect not only current taste for nuttiness and sensory intensity, but also how today’s sensory intensity alters tomorrow’s marginal utility. These equalities therefore encode the intertemporal wedge introduced by past consumption directly into the price system.

Finally, note that the mechanism tested here differs conceptually from inventory-driven persistence (Hendel and Nevo, 2006): stockpiling generates serial correlation in purchases through intertemporal substitution and storage under time-separable preferences. Contrastingly, persistence in my framework operates through state dependence in utility over characteristics, so current consumption of habit-forming attributes carries a continuation value by shifting future marginal valuations. In purchase data, however, the two mechanisms need not be empirically separable without additional structure linking purchases to consumption.

Testing consistency reduces to an empirical search for shadow prices and a discount factor satisfying (B1)–(B3). The system is nonlinear jointly in shadow prices and β , but becomes linear conditional on β . For any fixed discount factor, feasibility can therefore be assessed via a linear programme. Repeating this feasibility check over a grid of candidate discount factors yields a computationally straightforward implementation strategy.

A practical complication arises from condition (B1). In its raw form, cyclical monotonicity requires the inequalities to hold for all finite ordered cycles of observations, which quickly becomes computationally

burdensome as T grows. Online Appendix G derives an equivalent linear-programming formulation based on Afriat inequalities, replacing this cycle condition with a quadratic number of pairwise constraints in T .

Interpreting the Afriat test requires care. A positive result establishes existence: there exists *some* concave utility function and discount factor consistent with the data. The representation, however, is not unique. Distinct utility functions—beyond simple monotone transformations—may rationalise the same dataset. Moreover, rationalisability depends on the level of temporal and product aggregation; for example, time aggregation may smooth consumption in a way that mimics habit persistence.

A negative result is likewise not diagnostic of the precise source of failure. Rejection may reflect habit persistence extending beyond one lag, non-concavities in preferences, misspecification of the characteristics technology, or an incorrect partition of \mathbf{A} into contemporaneous and habit-forming components (for instance, treating “taste” preferences like nuttiness as static when the data in fact suggest that exposure today shifts future marginal valuations). The Afriat test should therefore be interpreted as a sharp but reduced-form diagnostic of dynamic rationalisability, rather than as a structural identification device.

2.4 A necessary rank condition for consistency

Although Theorem 2.1 provides a complete characterisation, its direct implementation can become computationally burdensome as the dimension of goods, characteristics, or time increases. This subsection therefore derives a low-dimensional diagnostic implied by consistency: a necessary rank condition that can be checked period by period. Failure of the condition immediately falsifies the model, while satisfaction is necessary but not sufficient.

The restriction is *structural* in the sense that it concerns only whether the maintained goods-to-characteristics technology can in principle support *any* shadow-price representation of observed prices. If the condition fails, no candidate preferences—static or dynamic—can rationalise the data because the implied shadow-price system does not exist.

Let $\boldsymbol{\rho}_t^+$ denote the $K_t^+ \leq K$ -dimensional vector of discounted prices of goods consumed in strictly positive quantities in period t . Let \mathbf{B}_t denote the $J \times K_t^+$ submatrix of \mathbf{A} collecting the contemporaneous characteristics of those goods, and let \mathbf{B}_t^a denote the $J_2 \times K_t^+$ submatrix collecting the corresponding habit-forming characteristics. For interior dates, Theorem 2.1 implies that

$$\boldsymbol{\rho}_t^+ = \mathbf{B}_t' \boldsymbol{\pi}_t^0 + (\mathbf{B}_t^a)' \boldsymbol{\pi}_{t+1}^1 = \left[\mathbf{B}_t' \mid (\mathbf{B}_t^a)' \right] \begin{bmatrix} \boldsymbol{\pi}_t^0 \\ \boldsymbol{\pi}_{t+1}^1 \end{bmatrix} =: \tilde{\mathbf{B}}_t \begin{bmatrix} \boldsymbol{\pi}_t^0 \\ \boldsymbol{\pi}_{t+1}^1 \end{bmatrix}, \quad (5)$$

where $\tilde{\mathbf{B}}_t$ is the $K_t^+ \times (J + J_2)$ augmented technology matrix. Because \mathbf{B}_t^a is formed by rows of \mathbf{B}_t , the augmented matrix $\tilde{\mathbf{B}}_t$ has the same column space as \mathbf{B}_t' , so habit formation tilts shadow prices within the same J -dimensional price manifold rather than expanding it. At a genuine terminal date, the same pricing equation holds with $\boldsymbol{\pi}_{T+1}^1 = \mathbf{0}$; the interior-date formulation is the one directly relevant for the empirical implementation, which does not treat the final observed purchase period as terminal.

Consistency therefore requires the observed price vector $\boldsymbol{\rho}_t^+$ to lie in the column space of $\tilde{\mathbf{B}}_t$, yielding the following necessary condition.

Proposition 2.1. If the data $\{\boldsymbol{\rho}_t, \mathbf{x}_t\}_{t=1}^T$ are consistent with the one-lag habits-over-characteristics

model for technology \mathbf{A} , then for every $t \leq T - 1$,

$$\text{rank}(\tilde{\mathbf{B}}_t \mid \boldsymbol{\rho}_t^+) = \text{rank}(\tilde{\mathbf{B}}_t) \leq \min\{K_t^+, J\}, \quad (\text{NC})$$

where \mid denotes horizontal concatenation.

Violation of (NC) at any single date is sufficient to reject the model, providing a sharp, low-dimensional falsification criterion that can be evaluated period by period. The restriction coincides with the necessary spanning condition under intertemporal separability in Blow et al. (2008): habit formation changes the *level* of shadow prices but does not expand the price manifold implied by the mapping from goods to characteristics. As such, (NC) is a structural constraint driven by the geometry of \mathbf{A} rather than by the curvature or stability of preferences. Because (NC) is only necessary, however, it does not guarantee consistency: even when (NC) holds, the inequalities for goods consumed at zero quantities may still be infeasible. Nevertheless, (NC) substantially reduces the feasible set and provides a fast diagnostic for empirical implementation.

Beyond its falsification role, the rank restriction also has limited identification content. For any t , (5) requires the observed price vector $\boldsymbol{\rho}_t^+$ to lie in the column space of $\tilde{\mathbf{B}}_t$. Because \mathbf{B}_t^a is formed by rows of \mathbf{B}_t , the augmented matrix $\tilde{\mathbf{B}}_t$ has the same column space as \mathbf{B}_t' , so the feasible set of price vectors is at most J -dimensional. When habit formation is present ($J_2 > 0$), the mapping from $(\boldsymbol{\pi}_t^0, \boldsymbol{\pi}_{t+1}^1)$ to $\boldsymbol{\rho}_t^+$ is therefore not one-to-one: reallocating shadow value between contemporaneous and lag components along the habit-forming directions can leave $\boldsymbol{\rho}_t^+$ unchanged. Viewed through (5) alone, the decomposition into contemporaneous and habit components is thus generically set-identified. In the running cereal example, prices may identify the overall shadow value of “nuttness” and “sensory intensity,” but not separately how much of the value of sensory intensity reflects current taste versus its continuation value through habits.

The full characterisation in Theorem 2.1 can nevertheless sharpen this conclusion. Condition (B1) links the stacked discounted shadow prices $\tilde{\boldsymbol{\pi}}_t$ across dates through cyclical monotonicity evaluated at the observed augmented bundles $\tilde{\mathbf{z}}_t$, and therefore uses quantity variation as well as prices. These additional restrictions can shrink the admissible set of decompositions relative to the price-span argument embodied in (NC). For example, the data may satisfy Theorem 2.1 under the dynamic model while violating the static restriction $\boldsymbol{\pi}_t^1 \equiv \mathbf{0}$, in which case the zero-habit specification is excluded from the identified set. Absent further structure, however, Theorem 2.1 still need not point-identify, or sign, the habit component separately from the contemporaneous one: (B1) disciplines the joint evolution of the stacked vectors $\tilde{\boldsymbol{\pi}}_t$, but does not in general undo the non-uniqueness in the decomposition induced by (5).

2.5 Missing prices

I now relax the assumption that prices are observed for all market goods. In many applications, prices are recorded only for goods that are actually purchased, giving rise to a missing price problem.⁵ Missing prices complicate RP analysis because imputing unobserved prices requires auxiliary assumptions. An alternative is to treat missing prices as unknowns and ask whether there exist values that render the data rationalisable. This existence approach should be understood as a partial-identification device: without further restrictions, one can always rationalise non-purchases by assigning prohibitively high unobserved prices, so the goal is

⁵Throughout, I assume that the technology matrix \mathbf{A} is known and time-invariant. This reflects settings where characteristics can be directly observed or constructed (e.g., nutritional content, design features, emissions ratings), even when market prices are only recorded for purchased items; in practice, missing prices are far more common than missing characteristics.

to characterise when the observed prices alone already force a violation (or allow rationalisation) under the maintained technology. I now formalise this approach.

Let $\boldsymbol{\rho}_t^+$ denote the $K_t^+ \leq K$ sub-vector of period- t discounted prices for goods with strictly positive demand, and let \mathbf{B}_t and \mathbf{B}_t^a denote the corresponding $J \times K_t^+$ and $J_2 \times K_t^+$ sub-matrices of \mathbf{A} and \mathbf{A}^a . Let $\boldsymbol{\rho}_t^0$, \mathbf{B}_t^0 , and $\mathbf{B}_t^{a,0}$ denote the complementary sub-vectors and sub-matrices associated with zero demand. The full discounted price vector is $\boldsymbol{\rho}_t = (\boldsymbol{\rho}_t^+, \boldsymbol{\rho}_t^0)$. I can then state an Afriat-type characterisation for the habits-over-characteristics model with missing prices.

Theorem 2.2. The following statements are equivalent:

- (A⁺) There exist missing prices $\{\boldsymbol{\rho}_t^0\}_{t=1}^T$ such that the completed price-quantity data are consistent with the model for technology \mathbf{A} .
- (B⁺) There exist shadow discounted prices $\{\boldsymbol{\pi}_t^0\}_{t=1}^T$, $\{\boldsymbol{\pi}_t^1\}_{t=1}^T$ and a discount factor $\beta \in (0, 1]$ such that,

$$0 \leq \sum_{m=1}^M \tilde{\boldsymbol{\pi}}_{t_m}' (\tilde{\mathbf{z}}_{t_{m+1}} - \tilde{\mathbf{z}}_{t_m}) \quad \forall M \geq 2,$$

$$\forall (t_1, \dots, t_M) \in \{1, \dots, T\}^M, t_{M+1} = t_1 \quad (\text{B1}^+)$$

$$\boldsymbol{\rho}_t^+ = \mathbf{B}_t' \boldsymbol{\pi}_t^0 + (\mathbf{B}_t^a)' \boldsymbol{\pi}_{t+1}^1 \quad \forall t \in \{1, \dots, T\} \quad (\text{B2}^+)$$

where $\tilde{\boldsymbol{\pi}}_t := \frac{1}{\beta^{t-1}} [\boldsymbol{\pi}_t^{0'}, \boldsymbol{\pi}_t^{1'}]'$, with $\boldsymbol{\pi}_{T+1}^1 \equiv \mathbf{0}$.

Proof: See Appendix A. \square

Relative to Theorem 2.1, the missing-price test conditions only on purchased-good prices and is therefore weaker. In empirical settings where the initial stock is unobserved and the observed sample window is not treated as the consumer's terminal horizon, the implementable counterpart is the observable interior-date analogue of Theorem 2.2, which evaluates the restrictions on dates for which both the lagged bundle and the one-step-ahead continuation term are observed.

3 Corollaries and nesting results

Before turning to the empirical application, I record two immediate corollaries that situate the habits-over-characteristics model within the RP literature. First, when characteristics coincide with market goods, the framework collapses to the habits-over-goods model of Crawford (2010). Second, under intertemporal separability and exponential discounting, it reduces to the characteristics-based model of Blow et al. (2008). Both results follow directly from Theorem 2.1 once the model is specialised to the relevant limiting cases. Formal definitions and derivations for these special cases are provided in Appendix A.

3.1 Habits-over-goods as a special case

When characteristics coincide with market goods—that is, when $J = K$ and the technology matrix satisfies $\mathbf{A} = \mathbf{I}_J$ —the habits-over-characteristics framework reduces to a standard habits-over-goods model. In this case, the distinction between goods and characteristics disappears, and the intertemporal first-order conditions involve current and lagged consumption of habit-forming goods directly. To match Crawford

(2010), I maintain the additional assumption that all goods are consumed in strictly positive quantities. Throughout this subsection, I adopt the terminal convention $\boldsymbol{\rho}_{T+1}^{a,1} \equiv \mathbf{0}$, so the final-period habit-good restriction has no continuation term. I also normalise the marginal utility of lifetime wealth to one, $\lambda = 1$, without loss of generality.

Definition 3.1. The data $\{\boldsymbol{\rho}_t^c, \boldsymbol{\rho}_t^a; \mathbf{x}_t^c, \mathbf{x}_t^a\}_{t=1}^T$ are *consistent* with the one-lag habits-over-goods model if there exists a locally non-satiated, superdifferentiable, and concave utility function $u(\cdot)$ and a positive constant β such that for all $t \in \{1, \dots, T\}$:

$$\boldsymbol{\rho}_t^c = \beta^{t-1} \partial_{\mathbf{x}_t^c} u(\bar{\mathbf{x}}_t), \quad (6)$$

$$\boldsymbol{\rho}_t^a = \beta^{t-1} \partial_{\mathbf{x}_t^a} u(\bar{\mathbf{x}}_t) + \beta^t \partial_{\mathbf{x}_t^a} u(\bar{\mathbf{x}}_{t+1}), \quad (7)$$

where $\boldsymbol{\rho}_t^c$ and $\boldsymbol{\rho}_t^a$ denote present-value prices of non-habit-forming and habit-forming goods, respectively.

Given this definition, I obtain the following result, equivalent to that in Crawford (2010) under the same normalisation.⁶

Corollary 3.1. The following statements are equivalent:

(A') The data $\{\boldsymbol{\rho}_t^c, \boldsymbol{\rho}_t^a; \mathbf{x}_t^c, \mathbf{x}_t^a\}_{t=1}^T$ are consistent with the one-lag habits-over-goods model.

(B') There exist T shadow price vectors $\{\boldsymbol{\rho}_t^{a,0}\}_{t=1}^T$, T shadow price vectors $\{\boldsymbol{\rho}_t^{a,1}\}_{t=1}^T$, and a positive constant β such that:

$$0 \leq \sum_{m=1}^M \tilde{\boldsymbol{\rho}}'_{t_m} (\bar{\mathbf{x}}_{t_{m+1}} - \bar{\mathbf{x}}_{t_m}) \quad \forall M \geq 2, \forall (t_1, \dots, t_M) \in \{1, \dots, T\}^M, t_{M+1} = t_1, \quad (\text{B1}')$$

$$\rho_t^{a,k} = \mathbf{e}_k \boldsymbol{\rho}_t^{a,0} + \mathbf{e}_k \boldsymbol{\rho}_{t+1}^{a,1} \quad \text{if } x_t^k > 0, \quad \forall k, t \in \{1, \dots, T\}, \quad (\text{B3}')$$

where $\bar{\mathbf{x}}_t := (\mathbf{x}_t^c, \mathbf{x}_t^a, \mathbf{x}_{t-1}^a)'$, \mathbf{e}_k is the k -th standard basis vector, and $\tilde{\boldsymbol{\rho}}_t := \frac{1}{\beta^{t-1}} [\boldsymbol{\rho}_t^c, \boldsymbol{\rho}_t^{a,0}, \boldsymbol{\rho}_t^{a,1}]'$.

Proof: See Appendix A. \square

3.2 Intertemporally separable preferences over characteristics

If habit formation is absent so that $\mathbf{z}_t = \mathbf{z}_t^c$ for all t , the model reduces to a characteristics-based framework with intertemporally separable preferences. Unlike Blow et al. (2008), who remain agnostic about intertemporal allocation, my formulation embeds this static characteristics model within a lifecycle problem with exponential discounting.⁷ The presence of a single intertemporal budget constraint implies a single shadow value of lifetime wealth, so observed choices must satisfy both intratemporal utility maximisation

⁶Equivalence uses the normalisation of the marginal utility of lifetime wealth $\lambda = 1$, which can be imposed without loss of generality.

⁷By a lifecycle problem I mean that the consumer chooses the entire consumption path to maximise lifetime utility subject to a single present-value budget constraint under exponential discounting.

over characteristics and intertemporal optimality. Consistency with this lifecycle formulation is defined as follows. As in the rest of the paper, I normalise the marginal utility of lifetime wealth to one.

Definition 3.2. The data $\{\boldsymbol{\rho}_t; \mathbf{x}_t\}_{t=1}^T$ are *consistent* with intertemporally separable preferences over characteristics and a life-cycle model for given technology \mathbf{A} if there exists a locally non-satiated, superdifferentiable, and concave utility function $u(\cdot)$ such that, for all $t \in \{1, \dots, T\}$,

$$\boldsymbol{\rho}_t \geq \mathbf{A}' \boldsymbol{\pi}_t, \quad (8)$$

with equality for all k such that $x_t^k > 0$, where $\mathbf{z}_t = \mathbf{A}\mathbf{x}_t$, $\boldsymbol{\rho}_t$ denotes the vector of present-value prices, and $\boldsymbol{\pi}_t := \partial_{\mathbf{z}_t} u(\mathbf{z}_t)$.

This recovers the hedonic pricing equation of Gorman (1956) and the first-order condition of Blow et al. (2008) when the static characteristics model is embedded in a lifecycle framework. From this, I obtain the following characterisation.

Corollary 3.2. The following are equivalent:

(A'') The data $\{\boldsymbol{\rho}_t; \mathbf{x}_t\}_{t=1}^T$ are consistent with intertemporally separable preferences over characteristics and a life-cycle model.

(B'') There exist T shadow price vectors $\{\boldsymbol{\pi}_t\}_{t=1}^T$ such that:

$$0 \leq \sum_{m=1}^M \boldsymbol{\pi}'_{t_m} (\mathbf{z}_{t_{m+1}} - \mathbf{z}_{t_m}) \quad \forall M \geq 2, \forall (t_1, \dots, t_M) \in \{1, \dots, T\}^M, t_{M+1} = t_1, \quad (\text{B1}'')$$

$$\rho_t^k \geq \mathbf{a}'_k \boldsymbol{\pi}_t \quad \forall k, t \in \{1, \dots, T\}, \quad (\text{B2}'')$$

$$\rho_t^k = \mathbf{a}'_k \boldsymbol{\pi}_t \quad \text{if } x_t^k > 0, \quad \forall k, t \in \{1, \dots, T\}, \quad (\text{B3}'')$$

where \mathbf{a}_k is the k th column of \mathbf{A} .

Proof: See Appendix A. \square

Together, these corollaries clarify the paper's value added. The framework unifies existing RP results in goods space and in characteristics space within a single dynamic hedonic model, and it shows how the interpretation of prices changes once habit formation over attributes is allowed. Relative to intertemporal separability, the dynamic model introduces continuation values into the shadow-price system and delivers diagnostics—via feasibility of (B1)–(B3) and the rank condition (NC)—that separate failures of the hedonic price representation from failures of dynamically coherent preferences.

4 Testing dynamic valuation in cereal purchases

This section applies the framework to household-level cereal purchases as a diagnostic of when characteristic-based valuation is admissible in scanner data, and when it requires a dynamic rather than static interpretation. Using these data, I show that most differences in raw rationalisability across models are driven by the

structural restrictions imposed by the hedonic representation of observed prices, while allowing for habits systematically improves *behavioural* coherence conditional on those constraints. The goal of the application is not to estimate a cereal-demand model, but to illustrate the framework in a realistic scanner setting where only purchases are observed and the mapping from purchases to underlying consumption states must be treated as a maintained approximation rather than as a directly observed object.

The empirical exercise proceeds sequentially. First, I ask a structural question: do observed within-period price vectors admit a hedonic representation? Equivalently, do they satisfy the equalities implied by the goods-to-characteristics technology (i.e., (B2⁺) in Theorem 2.2), so that characteristic shadow prices are well defined? These equalities act as a gatekeeper: if they fail, the behavioural test is not economically meaningful because the relevant shadow prices do not exist. Second, conditional on structural feasibility, I ask a behavioural question: do observed purchase sequences satisfy the dynamic RP inequalities (i.e., (B1⁺)) when evaluated at the implied shadow prices? In the data, these inequalities are applied to purchase-period bundles, so the resulting exercise should be read as a diagnostic of dynamic coherence under the maintained approximation that purchases track the relevant underlying consumption state at the chosen aggregation.

The sequential decomposition yields two findings. First, most variation in raw pass rates when moving from goods space to characteristics space reflects the hedonic price-system restrictions rather than differences in behavioural fit; empirically, however, restoring hedonic consistency typically requires only modest price adjustments. Second, conditional on satisfying those structural restrictions, allowing for habits systematically improves behavioural coherence. These results should therefore be read as evidence about where static and dynamic valuation place empirical discipline in scanner data.

Market goods are UPC-level products and characteristics are their nutritional content and a small set of descriptive indicators. The key empirical challenge is that prices are observed only for purchased items. Accordingly, I implement the missing-price test of Section 2.5.⁸ In this setting, behavioural discipline is inherently limited by sparse price support and modest intertemporal budget variation, so the diagnostic provides a lower bound on the model’s behavioural content. In environments with richer price and quantity support, the same framework may generate sharper behavioural restrictions.

4.1 Data

I use household-level scanner data on cold cereal purchases from the IRI Academic Datasets’ *BehaviorScan* panel. The panel covers two U.S. markets (Pittsfield, MA, and Eau Claire, WI) over 2010–2011, with purchases recorded at checkout via household ID cards and Universal Product Codes (UPCs). I restrict attention to static households that participate in all 12 months of a calendar year, so recruitment and attrition occur only at year-end.

I compute household-specific time periods to accommodate heterogeneity in purchase frequencies. For household i , let S_i denote the span from first to last purchase and G_i the longest interpurchase gap (including endpoints). I set $T_i = \lfloor S_i/G_i \rfloor$ and partition S_i into T_i equal-length bins, which guarantees at least one purchase in each period by construction. Households with $T_i < 3$ are excluded to ensure at least two observed transitions in the one-lag model. This aggregation is conservative for a one-lag specification: it ensures that the lagged bundle is observed rather than imputed, hence dynamic restrictions are evaluated on realised purchase transitions. After additionally dropping households with purchases lacking characteristics

⁸Imputing prices for unpurchased goods is feasible (e.g., using regional price indices), but it introduces auxiliary assumptions that can blur whether failures reflect preferences or the imputation procedure.

information (described below), the analysis sample contains $N = 2,282$ households.⁹ Table 1 summarises the resulting panel structure and the scale and variety of the implied choice problems.

Table 1: Sample and household-level panel structure

Statistic	Mean	Median
Households (N)	2,282	
Months covered	2010–2011 (24 months)	
Distinct products (K)	801	
Characteristics (J)	23	
Time periods per household (T_i)	7.0	6.0
Period length (days)	106	98
Units purchased per household-period	7.9	6.0
Total units purchased (24 months)	55.1	42.0
Expenditure per household-period (\$)	22.1	19.5
Total expenditure (24 months, \$)	154.2	116.5
Distinct products purchased (24 months)	32.3	25.0

Notes: Statistics are computed across $N = 2,282$ households. T_i denotes the number of constructed time periods for household i , defined as $T_i = \lfloor S_i/G_i \rfloor$, where S_i is the span between first and last purchase and G_i the longest interpurchase gap. Units are package counts; expenditures and package counts are aggregated within household-periods and totals are over 24 months.

Households purchase on average 7.9 units per period (median 6.0). Combined with a median period length of 98 days, this pattern is consistent with ongoing consumption rather than extreme purchase spikes. While promotions may induce some stockpiling, explaining the observed interpurchase gaps purely through inventory accumulation would require implausibly large stock build-ups relative to total quantities purchased. At the same time, within-period baskets remain narrow in UPC variety: the median household-period contains only two distinct purchased products, even though households typically buy several units. I therefore interpret each household-period as a meaningful dynamic choice observation, but not as a literal measure of contemporaneous nutrient intake: it accommodates idiosyncratic shopping frequencies while preserving the temporal structure required for testing whether aggregated purchase bundles behave in a manner consistent with the dynamic RP restrictions.¹⁰

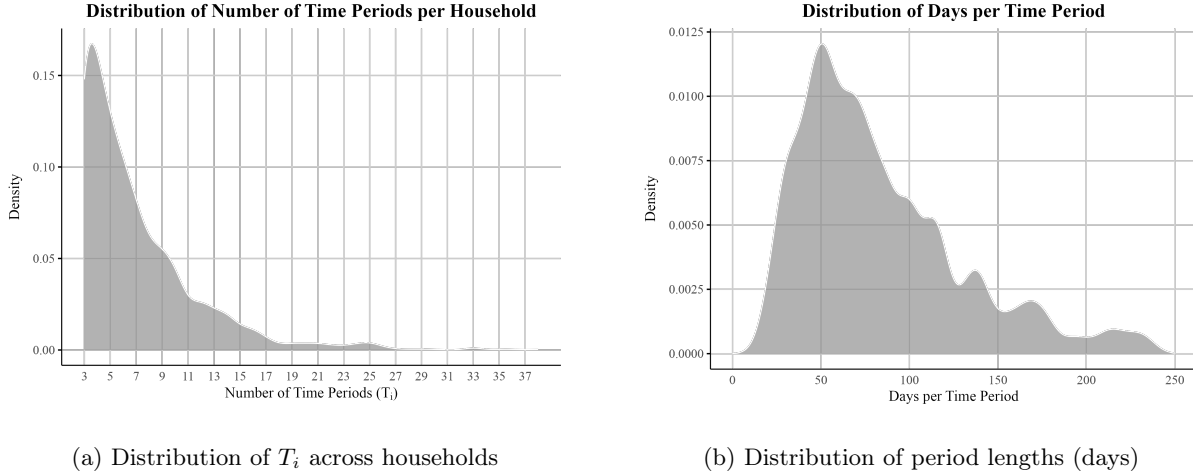
Figure 1 shows substantial heterogeneity in shopping frequency (median $T_i = 6$; median period length 98 days). Within each period I aggregate UPC-level quantities and nominal expenditures and compute unit values as expenditure-to-quantity ratios.¹¹ Because unpurchased goods have no observed unit values, the test treats missing prices as unknowns and searches for completions consistent with rationalisability. Reported pass rates are therefore upper bounds: a household that fails cannot be rescued by any price imputation, while a household that passes does so under at least one completion.

⁹Appendix C documents the full sequence of sample construction and reports balance tests comparing households excluded due to missing characteristics data with the final analysis sample. Differences are modest in magnitude, suggesting limited scope for selection on observables.

¹⁰Similar coarse aggregation is common in empirical work on dynamic demand when purchase occasions are intermittent (e.g., Crawford (2010) uses quarterly panels in an application to tobacco). I obtain qualitatively similar cross-model comparisons under a common monthly aggregation (i.e., $T_i = 24$ for all i), albeit in a much smaller effective sample because zero-purchase months become pervasive.

¹¹Expenditures are recorded in nominal dollars and converted to present-value terms for the lifecycle formulation using a monthly interest-rate series (30-Year Fixed Rate Mortgage Average, FRED (Freddie Mac, 2025)); results are unchanged under nominal values given the short sample window.

Figure 1: Distribution of constructed time periods and period lengths



Notes: Left panel shows the distribution of T_i across households ($N = 2,282$). Right panel shows the distribution of household-level average period length in days. $T_i = \lfloor S_i/G_i \rfloor$, where S_i is the span between first and last purchase and G_i the longest interpurchase gap.

The model is defined over characteristics rather than market goods. For each UPC I construct a vector of nutritional and descriptive attributes by merging the IRI product file to the (now-defunct) NuVal shelf-labelling database and, where NuVal is missing, supplementing with data from the FatSecret Platform API, following the approach in Barahona et al. (2023). I standardise all nutrients to a 100g basis so that continuous characteristics are comparable across products. Since scanner quantities are recorded in package counts, the resulting characteristics matrix should be interpreted as a linear summary of purchased product attributes. To handle IRI’s placeholder codes for some private-label items, I first map System-88 pseudo-UPCs to their corresponding real UPCs prior to merging. The merged characteristics cover over 97% of purchase-weighted observations; I drop unmatched purchase records and then drop households with any remaining unmatched purchases, so every household in the analysis sample has a complete characteristics mapping for all of its observed purchases.

I construct $J = 23$ characteristics. Eight are continuous Nutrition Facts Panel measures (calories, carbohydrates, total fat, saturated fat, fibre, protein, sodium, and sugar). The remaining 15 are binary indicators capturing salient ingredients and descriptors (10 indicators) and the five most prevalent brands (Kellogg’s, General Mills, Post, Quaker, and Kashi). Table 2 reports the binary definitions. In the baseline specification I treat sugar and sodium as habit-forming characteristics ($J_2 = 2$) and take the remaining characteristics as non-habit-forming, motivated by evidence that sugar and salt may activate reward pathways in ways analogous to addictive substances (Avena et al., 2008; Cocores and Gold, 2009). Section 4.2.1 reports robustness to alternative partitions. Additional descriptive evidence on purchase intensity, brand concentration, and the distributions of prices and characteristics is reported in Appendix C.

Table 2: Binary product characteristics used in the hedonic representation

Binary characteristic	Construction
whole grain	contains ‘whole grain’ or ‘wholegrain’ in ingredient list
organic	contains the term ‘organic’ at least once in ingredient list
oat-based	contains ‘oat’ as part of the first listed ingredient
granola	contains ‘granola’ in product description
health halos	contains ‘source of’, ‘natural’, ‘low calorie’, or ‘low fat’ in product description
fruity	contains ‘fruit’ in the product description
nutty	contains ‘nut’ in product description
chocolatey	contains ‘chocolate’ in product description
honey flavour	contains ‘honey’ in product description
gluten-free	contains ‘gluten free’ in product description
Kellogg’s	cereal brand is Kellogg’s
General Mills	cereal brand is General Mills
Post	cereal brand is Post
Quaker	cereal brand is Quaker
Kashi	cereal brand is Kashi

Notes: Characteristics are defined at the UPC level using ingredient lists and product descriptions from the merged IRI–NuVal–FatSecret dataset. Indicators equal one if the stated textual condition is satisfied. Brand indicators correspond to the five most prevalent national brands in the sample. Continuous nutritional characteristics (calories, carbohydrates, total fat, saturated fat, fibre, protein, sodium, and sugar, standardised per 100g) are defined separately in Appendix C.

4.2 Results

4.2.1 Rationalisability scores

Two patterns emerge in the raw rationalisability outcomes: allowing for habits increases pass rates within characteristics space, and goods-based representations exhibit substantially higher pass rates. As the sequential logic of the paper makes clear, however, these raw differences conflate structural and behavioural components. Section 4.2.2 decomposes these margins and quantifies the severity of violations.

The test is implemented at the household level, allowing full heterogeneity in pass/fail outcomes and in the shadow-price structure of rationalisable households. Because habits enter with a one-period lag, the empirical exercise uses the observable interior-date analogue of Theorem 2.2. Structural equalities are imposed on $t = 2, \dots, T - 1$, since period $t = 1$ depends on the unobserved initial stock and I do not treat the last observed purchase period as the consumer’s terminal horizon. The behavioural Afriat inequalities are then evaluated on the retained dates $t = 2, \dots, T$, so the final observed period continues to discipline the shadow-price sequence. Unless otherwise specified, I impose a lifecycle model with a single intertemporal budget constraint, so the marginal utility of income is constant across periods.

Under the baseline habits-over-characteristics specification, 1,248 of 2,282 households (54.69%) satisfy the test. Varying the set of habit-forming characteristics has essentially no effect on classification: restricting habits to sugar alone (54.60%), sodium alone (54.65%), or allowing all 23 characteristics to be habit-forming (54.69%) changes the pass rate by at most 0.09 percentage points and reclassifies no more than two households. In this application, the data therefore do not isolate sugar or sodium as uniquely responsible for the dynamic improvement; rather, the main empirical action comes from the maintained hedonic representation, with habits mattering at the behavioural margin once that representation is imposed. Table 3 reports the full set of specifications.

Table 3: Household-level pass rates: habits-over-characteristics specifications

Model	Description	Lifecycle model	Pass rate
Habits-over-characteristics	$J_2 = 2$ (sugar and sodium)	Yes	54.69%
Habits-over-sugar	$J_2 = 1$ (sugar)	Yes	54.60%
Habits-over-sodium	$J_2 = 1$ (sodium)	Yes	54.65%
Habits-over-all-characteristics	$J_2 = J = 23$	Yes	54.69%

Notes: A household passes if there exist characteristic shadow prices satisfying the interior-date structural equalities (B2⁺) in Theorem 2.2 on $t = 2, \dots, T - 1$ and the corresponding dynamic RP conditions (B1⁺) on $t = 2, \dots, T$; by Theorem 2.2, this is equivalent to the existence of a valid completion of missing prices in the background. Pass rates are computed over $N = 2,282$ households. $J = 23$ denotes the total number of characteristics and J_2 the number assumed to be habit-forming. All specifications impose a lifecycle model with a single intertemporal budget constraint.

Raw pass rates differ sharply across representations (Table 4). Moving from characteristics to goods sharply increases rationalisability from 54.7% to 99.6% under dynamic preferences. This large difference is driven by the much greater dimensional flexibility of goods-based representations, which impose no cross-good price restrictions. In this application, the structural restrictions bite through the geometry of the active purchased bundle rather than through broad within-period variety. Many of the apparent “failures” in characteristics space nevertheless correspond to economically small deviations from those hedonic restrictions. Section 4.2.2 makes this precise by separating structural and behavioural sources of empirical discipline and quantifying the magnitude of violations.

By contrast, removing habits within characteristics space reduces the pass rate from 54.7% to 52.4%. While this difference in levels is modest, the paired comparisons below show that the reclassification is entirely directional, with households failing under static preferences but passing once dynamics are introduced. I show later that allowing for habits also reduces the severity of behavioural violations on average, but the empirical gains are concentrated rather than universal.

Table 4: Household-level pass rates across characteristics and goods representations

Model	Description	Lifecycle model	Pass rate
Habits-over-characteristics	$J_2 = 2$ (sugar and sodium)	Yes	54.69%
Characteristics (no habits)	$J_2 = 0$	Yes	52.37%
Habits-over-all-goods	$K = J = J_2$, $\mathbf{A} = \text{identity}(K)$	Yes	99.56%
Goods (no habits)	$K = J = J_1$, $\mathbf{A} = \text{identity}(K)$	Yes	92.59%
Goods (no habits) (GARP)	$K = J = J_1$, $\mathbf{A} = \text{identity}(K)$	No	99.69%

Notes: A household passes if there exist characteristic shadow prices satisfying the interior-date structural equalities (B2⁺) in Theorem 2.2 on $t = 2, \dots, T - 1$ and the corresponding dynamic RP conditions (B1⁺) on $t = 2, \dots, T$; by Theorem 2.2, this is equivalent to the existence of a valid completion of missing prices in the background. Pass rates are computed over $N = 2,282$ households. In characteristics models, $J = 23$ and J_2 denotes the number of habit-forming characteristics. In goods models, K denotes the number of goods and $\mathbf{A} = \text{identity}(K)$ implies no cross-good price restrictions. “Lifecycle model” indicates whether a single intertemporal budget constraint is imposed.

The reclassification pattern is strongly directional. Removing habits from the baseline characteristics model reclassifies 53 households, all of whom fail under static preferences but pass once intertemporal dependence is introduced (p-value $< 10^{-15}$). These households provide the cleanest empirical illustration of the valuation wedge in Section 2.2: the maintained characteristics technology can support a hedonic interpretation, but only if current characteristics are allowed to carry continuation values. By contrast,

changing the allocation of habit-forming characteristics reclassifies at most three households and yields no statistically significant differences. Table 5 reports the full set of paired comparisons.

Comparisons with goods-based models reveal even larger directional differences. In these cases, most switching households fail the characteristics-based test but pass the corresponding goods-based alternative. As shown in the next subsection, this pattern reflects the much greater dimensional flexibility of the goods representation, which imposes no cross-good price restrictions, rather than tighter behavioural alignment.

Table 5: Paired model comparisons (exact McNemar tests)

Comparison (baseline vs alternative)	Pass ₀	Pass ₁	Δ (pp)	Switchers	p -value
Baseline: Habits-over-characteristics	54.69				
Habits-over-all-characteristics	54.69	54.69	0.00	0	1.000
Habits-over-sugar	54.69	54.60	0.09	2	0.500
Habits-over-sodium	54.69	54.65	0.04	1	1.000
Characteristics (no habits)	54.69	52.37	2.32	53	$< 10^{-15}$
Habits-over-all-goods	54.69	99.56	-44.87	1024	$< 10^{-15}$
Goods (no habits)	54.69	92.59	-37.91	1039	$< 10^{-15}$
Goods (no habits) (GARP)	54.69	99.69	-45.00	1031	$< 10^{-15}$

Notes: Pass₀ is the baseline pass rate (habits-over-characteristics with $J_2 = 2$), Pass₁ is the alternative model pass rate, and Δ (pp) reports Pass₀ minus Pass₁ in percentage points. “Switchers” is $n(1 \rightarrow 0) + n(0 \rightarrow 1)$, i.e., the number of households whose pass/fail classification differs across the paired models. p -values are from exact McNemar (binomial) tests based on the switchers.

Taken together, these results indicate that while the precise *allocation* of habits across characteristics plays little role in classification, introducing intertemporal dependence in characteristics space improves rationalisability for a limited but directionally important subset of households. In this application, the main empirical action comes from the maintained hedonic representation, while habits matter at the behavioural margin. The results are therefore best read as showing how the diagnostic separates failures of the hedonic price representation from failures of the static interpretation, rather than as isolating sugar or sodium as uniquely responsible for dynamic behaviour.

4.2.2 Structural and behavioural sources of empirical discipline

If the test is to serve as an interpretation check for recovered WTP objects, rejection cannot be the end of the analysis. A binary pass/fail outcome does not reveal whether the static approximation fails badly or only marginally, nor whether rejection reflects the hedonic price representation or intertemporal behaviour. I therefore separate rationalisability into structural and behavioural components and quantify violations on each margin using continuous discrepancy measures. The construction exploits the fact that representation theorems induce measures of rationality violations through the distance to their defining restrictions (e.g., Andrews (2026)).

In the hedonic setting, this distinction admits a natural decomposition. Rationalisability in the habits-over-characteristics model requires joint satisfaction of two conceptually distinct restrictions. *Structural equalities* (B2⁺) link observed prices to product characteristics through the hedonic technology and act as overidentifying restrictions on the admissible shadow-price representation. This is the gatekeeper stage of the empirical analysis. *Behavioural inequalities* (B1⁺) constrain intertemporal choice through shadow prices and the discount factor (Theorem 2.2). Separating these margins clarifies where empirical discipline originates in characteristics-based valuation and how it differs from more flexible goods-based representations.

The structural equalities require that, in each household-period, the observed price vector ρ_t^+ lies in the column space of the augmented characteristics matrix $\tilde{\mathbf{B}}_t := [\mathbf{B}'_t \mid (\mathbf{B}_t^a)']$. In principle this is a dimensionality-reduction restriction: prices must be representable through a lower-dimensional characteristics technology rather than arbitrary goods-specific shifters. In the scanner environment studied here, however, the active choice sets are typically narrow, so the empirical bite of the restriction depends on the realised rank of the purchased bundle. When the purchased-good matrix has full row rank, the column space of $\tilde{\mathbf{B}}_t$ spans all of $\mathbb{R}^{K_t^+}$ and the structural equalities are mechanically satisfied. Non-zero structural distances therefore arise precisely in those household-periods where the active bundle is rank deficient, either because the household buys too few linearly independent products or because the purchased UPCs have highly similar characteristic profiles.

This is empirically plausible in cereal data, where closely related UPCs often differ only in package size, branding, or minor formulation details and therefore carry very similar measured characteristic vectors. Even after conditioning on a rich set of nutritional and descriptive characteristics, scanner prices also reflect retailer pricing strategies, temporary promotions, and mark-ups driven by market power or shelf placement that do not correspond to attributes households consume. When the active bundle is low rank, those retailer-specific price components cannot be absorbed by the hedonic system and therefore appear as structural residuals.

To quantify the severity of these violations, I compute for each household-period the Euclidean distance

$$d_t = \|\rho_t^+ - \tilde{\mathbf{B}}_t \tilde{\mathbf{B}}_t^+ \rho_t^+\|,$$

where $\tilde{\mathbf{B}}_t^+$ denotes the Moore–Penrose pseudoinverse and $\tilde{\mathbf{B}}_t \tilde{\mathbf{B}}_t^+$ is the orthogonal projector onto the equality manifold implied by the hedonic representation. Equivalently, observed prices admit the orthogonal decomposition

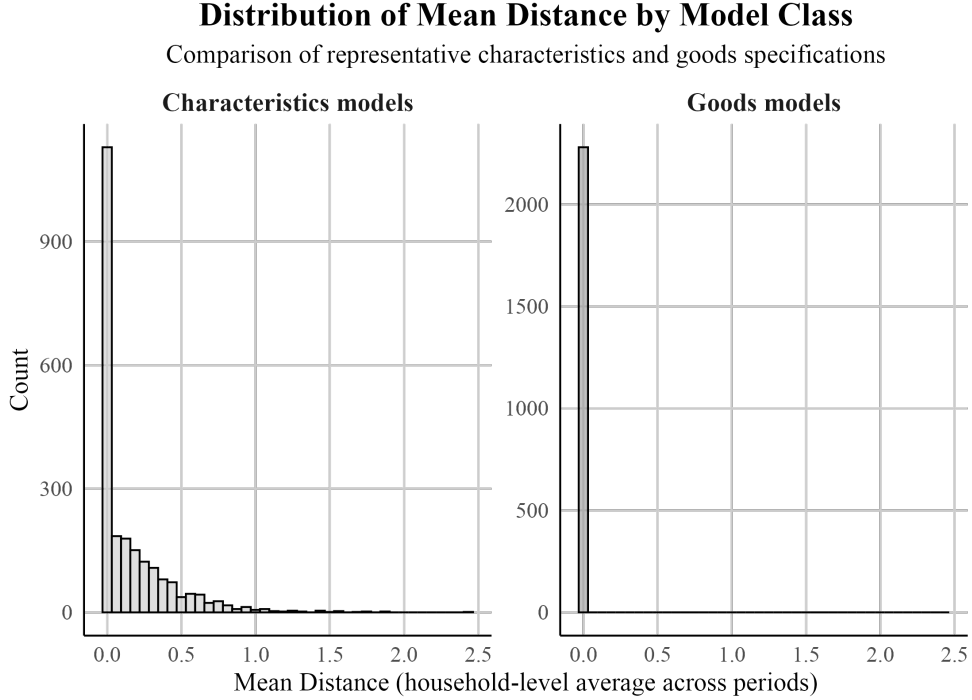
$$\rho_t^+ = \tilde{\mathbf{B}}_t \hat{\pi}_t + \mathbf{r}_t,$$

where $\hat{\pi}_t$ minimises $\|\rho_t^+ - \tilde{\mathbf{B}}_t \pi\|$ and \mathbf{r}_t is orthogonal to the column space of $\tilde{\mathbf{B}}_t$. The distance $d_t = \|\mathbf{r}_t\|$, measured in dollars, therefore captures the minimal joint price adjustment required for a hedonic price representation to exist. For example, $d_t = 1.0$ means that the minimum adjustment vector to prices has Euclidean norm 1 dollar; equivalently, the sum of squared price adjustments across the goods purchased that period is 1.

One response to violations of the hedonic equalities is to augment the technology with latent characteristics, as in Blow et al. (2008), thereby expanding the price manifold until the equalities hold. I do not pursue this route here. My objective is not to restore rationalisability by construction, but to quantify how demanding a given hedonic representation is in the data and to separate structural misspecification from behavioural inconsistency.

Figure 2 plots the distribution of mean distances across households for both characteristics- and goods-based specifications. The characteristics models exhibit a wide, right-skewed distribution centred above zero. In this application, those distances should be read as evidence that many active bundles fail to span their own purchased-good price space once prices are projected through the maintained characteristics technology. Even so, the implied unit-price violations are often modest in magnitude: most household-periods require a minimum joint price adjustment with Euclidean norm below 1 dollar to restore hedonic consistency. As a rough benchmark, average household cereal expenditure is \$22.1 per period. By contrast, goods-based models impose no cross-good price restrictions and therefore satisfy the structural equalities mechanically, yielding $d_t = 0$.

Figure 2: Mean distance-to-manifold



Notes: Histograms of the household-level mean of the structural distance measure across periods. Distances are measured in dollars and represent the minimal joint price adjustment required for the hedonic price equalities to hold. Left panel reports the characteristics specification; right panel the goods specification.


An important implication of this geometry is that all characteristics-based specifications impose identical structural restrictions. Habit-forming characteristics enter $\tilde{\mathbf{B}}_t$ only as duplicated rows of \mathbf{B}_t and therefore neither increase its rank nor enlarge the feasible price set. Structural restrictiveness is thus governed entirely by the hedonic representation itself, not by the allocation of habits across characteristics. In the present data, this means that structural rejection is best understood as a statement about the linear algebra of narrow purchased bundles—their rank and characteristic collinearity—rather than about the dynamic specification per se. Moreover, these structural restrictions operate at the level of prices and characteristics and are conceptually distinct from inventory dynamics: even if households smooth consumption through inventories, prices must still admit a hedonic representation for shadow prices to be well defined.

To measure behavioural restrictiveness, I adapt the Critical Cost Efficiency Index (CCEI) following Afriat (1973) and Varian (1990). The CCEI measures the smallest proportional relaxation of revealed affordability required for the behavioural inequalities in $(\mathbf{B}1^+)$ to admit a solution.¹² Economically, a CCEI of η indicates that the revealed-affordability comparisons in the lifecycle RP test need only be relaxed by at most $(1 - \eta) \times 100\%$ for the observed intertemporal choice path to become rationalisable. Values close to one therefore indicate that behaviour is nearly dynamically rational, while lower values signal more severe violations. The index is defined only for households whose prices lie exactly on the equality manifold, because

¹²The CCEI is interpreted here strictly as a measure of RP slack, following its original cost-efficiency interpretation in Afriat (1973). As emphasized by Echenique (2022), it should *not* be interpreted as a welfare loss or a measure of foregone surplus.

characteristic shadow prices—and hence the behavioural inequalities themselves—are well defined only when the hedonic equalities hold. In this sense, the CCEI plays an analogous role to a goodness-of-fit statistic for intertemporal choice: it quantifies how much slack must be introduced before the model can rationalise observed behaviour, holding the price system fixed.

Figure 3: Critical Cost Efficiency Index (CCEI)



Notes: Histograms of household-level CCEI values across model specifications. The CCEI is the smallest proportional relaxation of the behavioural RP inequalities required for feasibility; values closer to one indicate smaller violations. Computed only for households satisfying the structural price equalities. Vertical axis is on a log scale.

Figure 3 shows that behavioural violations are generally modest, but systematically smaller when habits are allowed. All characteristics-based specifications exhibit substantial mass near unity, indicating that once the hedonic structure is satisfied, only limited perturbations are needed to rationalise behaviour. However, the static characteristics model without habits displays a thicker lower tail, reflecting greater intertemporal inconsistency even when structural feasibility holds. A parallel pattern arises in the goods domain: the habits-over-goods specification yields CCEIs tightly concentrated near one, while the static goods model exhibits a wider distribution with more mass below 0.95. These patterns indicate that habit formation strengthens behavioural discipline by reducing intertemporal reversals.

This behavioural improvement matters for interpretation. For households that satisfy the dynamic characteristics model but fail the static counterpart, the data reject the restriction that characteristic shadow prices can be read purely as contemporaneous marginal valuations. In those cases, a valuation exercise that treats sugar or sodium prices as static WTP imposes an additional timing restriction not supported by the observed purchase path. The empirical test therefore identifies when the dynamic wedge characterised in Section 2.2 is supported.

The above distance measures show where and how severely the model fails, but they do not by themselves show whether the restrictions are demanding in this scanner environment. This is the fit-versus-restrictiveness concern emphasised by Selten (1991): a model may rationalise many datasets because it captures structure, or because nearby alternatives would also pass. I therefore complement the decomposition with a simulation benchmark in the spirit of Fudenberg et al. (2023), which evaluates how unusually well the model fits the observed data relative to nearby feasible perturbations of the scanner environment. The benchmark preserves each household’s zero pattern and total expenditure while allowing prices and quantities to vary locally, thereby inducing a comparison distribution over structural and behavioural discrepancy measures. Full details are provided in Appendix E.

On the structural margin, the benchmark confirms that the hedonic equalities impose substantial empirical discipline. Observed scanner prices are, on average, closer to the equality manifold than locally perturbed price systems, yet are rarely extreme outliers relative to the induced comparison distribution. This indicates that the hedonic restrictions are demanding and not trivially satisfied. By contrast, goods-based models impose no cross-good price restrictions and therefore cannot fail on the structural margin, accounting for their much higher raw pass rates.

On the behavioural margin, CCEI values are extremely close to one in both the observed and locally perturbed data, leaving little scope for behaviour to appear unusually efficient in a quantile sense. This reflects the empirical environment—sparse active choice sets and limited intertemporal budget variation—rather than a lack of behavioural content in the model.

Taken together, the price geometry, behavioural slack measures, and simulation benchmark clarify the interpretation of raw pass rates. Structural restrictions embedded in the hedonic representation largely determine differences across goods and characteristics models, but in this application they do so through the rank properties of the active purchased bundles rather than through a broad market-level dimension count alone. Conditional on those restrictions, allowing for habit formation systematically improves intertemporal coherence relative to static preferences. Goods-based models achieve high pass rates primarily because they impose little structural content, not because they deliver a tighter account of behaviour.

4.2.3 Discount factors

The RP conditions are evaluated conditional on the discount factor β . Appendix D reports the fraction of rationalisable households consistent with each value on a fine grid $\beta \in [0.95, 1]$. Acceptance probabilities are uniformly high across the grid under both habits-over-characteristics and habits-over-goods specifications, with no economically meaningful monotonic pattern.

These results indicate that, conditional on the hedonic price restrictions, the behavioural inequalities impose only weak discipline on intertemporal discounting in this environment. The identification sets for β are wide, cautioning against interpreting discount factors recovered from nonparametric dynamic RP tests as tightly identified structural parameters in scanner data.

4.2.4 Predictors of rationalisability

Pass/fail outcomes are primarily associated with the scale and complexity of the observed choice problem rather than with demographics per se. Table 6 reports average marginal effects from probit regressions of the pass indicator on household characteristics. Column (1) includes demographics only; Column (2) adds measures of purchasing intensity and variety.

In the demographic-only specification, households with children are 10.6 percentage points less likely to pass. Once purchasing controls are added, this effect attenuates and becomes statistically insignificant. By contrast, scale measures remain economically large and significant: an interquartile increase in two-year cereal expenditure is associated with roughly a 25 percentage point lower probability of passing, while a comparable increase in product variety reduces the probability by about 8 percentage points.

Taken together, these results suggest that household composition matters primarily through the scale and complexity of the observed choice problem. Larger and more diverse purchasing patterns generate a greater number of structural and behavioural constraints, increasing the likelihood that at least one is violated. For instance, when a household purchases products with sharply different nutritional profiles across periods, the implied cross-attribute shadow-price system must rationalise a wider range of price–quantity trade-offs than in households that repeatedly buy a narrow set of similar cereals. Demographics per se have limited explanatory power once this effective dimensionality is accounted for. The demographics-only specification has almost no explanatory power (McFadden pseudo- $R^2 = 0.009$), whereas the specification with purchasing controls reaches a more substantial 0.176.

Table 6: Determinants of rationalisability: probit average marginal effects

	<i>Dependent variable: Pr(Pass)</i>	
	(1) demographics only	(2) w/ purchasing controls
young child	−0.106 (0.084)	−0.081 (0.079)
children	−0.106** (0.033)	−0.012 (0.030)
age group HH head	0.004 (0.012)	−0.006 (0.011)
education HH head	−0.00004 (0.008)	0.001 (0.007)
Pittsfield	0.005 (0.021)	−0.001 (0.022)
high income HH [†]	−0.036 (0.025)	0.008 (0.023)
total units purchased		0.00002 (0.001)
total expenditure		−0.002*** (0.0003)
total unique products		−0.003** (0.001)
Pseudo- R^2 (McFadden)	0.009	0.176
Observations	2,187	2,187

*p<0.1; **p<0.05; ***p<0.01.

[†]High-income indicator equals one for pre-tax income \geq 75k.

Notes: Entries report average marginal effects from probit models of an indicator for passing the habits-over-characteristics test. Column (1) includes demographic controls only; column (2) additionally includes purchasing intensity measures. Standard errors in parentheses are based on the delta method. 95 observations omitted due to missing demographic data.

5 Conclusion

This paper provides a nonparametric foundation for valuing characteristics when those characteristics may be habit forming. Policies often target attributes such as sugar, sodium, alcohol, caffeine, or emissions, and applied work often interprets characteristic prices or demand coefficients as static WTP. I show that this interpretation is valid only under a restriction: current consumption of the characteristic must carry no continuation value.

The main result is an Afriat-type characterisation of dynamic hedonic valuation. Observed prices and choices admit a dynamic characteristics-based interpretation if and only if there exist characteristic shadow prices and a discount factor satisfying structural price equalities and dynamic behavioural inequalities. The structural equalities require observed goods prices to be representable through the maintained goods-to-characteristics technology. The behavioural inequalities require choices to be dynamically rationalisable at the implied shadow prices.

The framework clarifies why dimensionality reduction and dynamics play different roles. Moving from goods to characteristics yields parsimony but imposes geometric discipline on the admissible representation of prices. Accommodating dynamic preferences does not relax that discipline. Rather, it changes the behavioural restrictions imposed once characteristic shadow prices exist: choices that violate static separability may still be rationalisable when current characteristics are allowed to affect future marginal valuations. This distinction matters for policy because, in such cases, the relevant WTP object is no longer purely contemporaneous.

The cereal application illustrates this logic in a purchase-based scanner setting. Goods-based benchmarks pass at very high rates, while characteristics-based models fail more often because observed prices frequently violate the hedonic spanning restriction. The associated distance measures show, however, that many structural violations are economically modest. Conditional on structural admissibility, behaviour is close to dynamically rational, and allowing for habit formation improves behavioural coherence relative to static characteristics models for a subset of households.

The framework is best understood as a diagnostic of admissibility and interpretation rather than as a point estimator of welfare primitives. Rationalisability does not identify a unique dynamic preference representation, and the continuation component of characteristic prices is generally set identified without further structure. This limitation is also informative: it shows exactly which additional assumptions are required before static WTP, dynamic WTP, or policy counterfactuals can be given a unique welfare interpretation.

Several directions for future work follow directly. Richer data with denser price support, more frequent observation, and greater intertemporal budget variation would sharpen both structural and behavioural restrictions. On the modelling side, extensions to stochastic choice, richer unobserved heterogeneity, and environments in which utility depends on current goods while habits attach to lagged characteristics would broaden the scope of the test. The broader message is that WTP for product characteristics, whether recovered from prices or demand systems, is not automatically a static welfare object: this paper provides a disciplined framework for testing whether characteristic shadow values exist and whether their interpretation must account for continuation effects.

References

- Afriat, S. N. (1967). The Construction of Utility Functions from Expenditure Data. *International Economic Review*, 8(1):67–77.
- Afriat, S. N. (1973). On a system of inequalities in demand analysis: An extension of the classical method. *International Economic Review*, 14(2):460–472.
- Allcott, H., Gentzkow, M., and Song, L. (2022). Digital addiction. *American Economic Review*, 112(7):2424–2463.
- Andrews, I. (2026). Revealed rationality: Label-free regularization from representation theorems. First version: February 21, 2026. Working paper, Department of Economics, MIT and NBER.
- Avena, N. M., Rada, P., and Hoebel, B. G. (2008). Evidence for sugar addiction: Behavioral and neurochemical effects of intermittent, excessive sugar intake. *Neuroscience & Biobehavioral Reviews*, 32(1):20–39.
- Bajari, P. and Benkard, C. L. (2005). Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach. *Journal of Political Economy*, 113(6):1239–1276.
- Bajari, P. and Kahn, M. E. (2005). Estimating Housing Demand With an Application to Explaining Racial Segregation in Cities. *Journal of Business & Economic Statistics*, 23(1):20–33.
- Barahona, N., Otero, C., and Otero, S. (2023). Equilibrium effects of food labeling policies. *Econometrica*, 91(3):839–868.
- Becker, G. S., Grossman, M., and Murphy, K. M. (1994). An Empirical Analysis of Cigarette Addiction. *The American Economic Review*, 84(3):396–418.
- Becker, G. S. and Murphy, K. M. (1988). A Theory of Rational Addiction. *Journal of Political Economy*, 96(4):675–700.
- Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–890.
- Blow, L., Browning, M., and Crawford, I. (2008). Revealed preference analysis of characteristics models. *The Review of Economic Studies*, 75(2):371–389.
- Boyer, M. (1978). A Habit Forming Optimal Growth Model. *International Economic Review*, 19(3):585–609.
- Boyer, M. (1983). Rational demand and expenditures patterns under habit formation. *Journal of Economic Theory*, 31(1):27–53.
- Browning, M. (1989). A Nonparametric Test of the Life-Cycle Rational Expectations Hypothesis. *International Economic Review*, 30(4):979–992.
- Cocores, J. A. and Gold, M. S. (2009). The Salted Food Addiction Hypothesis may explain overeating and the obesity epidemic. *Medical Hypotheses*, 73(6):892–899.
- Crawford, I. (2010). Habits Revealed. *The Review of Economic Studies*, 77(4):1382–1402.

- 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.
- Diewert, W. E. (1973). Afriat and Revealed Preference Theory. *The Review of Economic Studies*, 40(3):419.
- Dubois, P., Griffith, R., and O'Connell, M. (2020). How Well Targeted Are Soda Taxes? *American Economic Review*, 110(11):3661–3704.
- Echenique, F. (2022). On the meaning of the critical cost efficiency index.
- Felippa, C. A. (2004). *Introduction to finite element methods*. University of Colorado.
- Freddie Mac (2025). 30-year fixed rate mortgage average in the united states (mortgage30us). FRED, Federal Reserve Bank of St. Louis. Retrieved August 16, 2025.
- Fudenberg, D., Gao, W., and Liang, A. (2023). How Flexible is that Functional Form? Quantifying the Restrictiveness of Theories. *SSRN Electronic Journal*.
- Gibbons, S. and Machin, S. (2003). Valuing English primary schools. *Journal of Urban Economics*, 53(2):197–219.
- Gorman, W. M. (1956). A Possible Procedure for Analysing Quality Differentials in the Egg Market. *The Review of Economic Studies*, 47(5):843–856.
- Greenstone, M. and Gallagher, J. (2008). Does hazardous waste matter? evidence from the housing market and the superfund program. *The Quarterly Journal of Economics*, 123(3):951–1003.
- Grossman, M., Chaloupka, F. J., and Sirtalan, I. (1998). An Empirical Analysis of Alcohol Addiction: Results from the Monitoring the Future Panels. *Economic Inquiry*, 36(1):39–48.
- Gruber, J. and Kőszegi, B. (2004). Tax incidence when individuals are time-inconsistent: the case of cigarette excise taxes. *Journal of Public Economics*, 88(9-10):1959–1987.
- Haeck, C., Lawson, N., and Poirier, K. (2022). Estimating consumer preferences for different beverages using the BLP approach. *Economics & Human Biology*, 46:101138.
- Heckman, J. and Scheinkman, J. (1987). The Importance of Bundling in a Gorman-Lancaster Model of Earnings. *The Review of Economic Studies*, 54(2):243–255.
- Hendel, I. and Nevo, A. (2006). Sales and consumer inventory. *The RAND Journal of Economics*, 37(3):543–561.
- Houthakker, H. S. (1950). Revealed Preference and the Utility Function. *Economica*, 17(66):159–174.
- Koob, G. F. and Moal, M. L. (1997). Drug Abuse: Hedonic Homeostatic Dysregulation. *Science*, 278(5335):52–58.
- Le Fur, E. and Outreville, J. F. (2022). The Willingness to Pay for Cider Products: Results of a Survey on Habits and Consumption Behavior. *Wine Business Journal*, 5(1).
- Nevo, A. (2001). Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica*, 69(2):307–342.

- Richards, T. J. and Patterson, P. M. (2006). Native American Obesity: An Economic Model of the "Thrifty Gene" Theory. *American Journal of Agricultural Economics*, 88(3):542–560.
- Rockafellar, R. T. (1970). *Convex Analysis*, volume 28 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ.
- Rosen, S. (1974). Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy*, 82(1):34–55.
- Samuelson, P. A. (1948). Consumption Theory in Terms of Revealed Preference. *Economica*, 15(60):243–253.
- Selten, R. (1991). Properties of a measure of predictive success. *Mathematical Social Sciences*, 21(2):153–167.
- Smith, V. K. and Desvousges, W. H. (1986). The Value of Avoiding a Lulu: Hazardous Waste Disposal Sites. *The Review of Economics and Statistics*, 68(2):293–299.
- Varian, H. R. (1982). The Nonparametric Approach to Demand Analysis. *Econometrica*, 50(4):945–973.
- Varian, H. R. (1990). Goodness-of-fit in optimizing models. *Journal of Econometrics*, 46(1-2):125–140.

Appendix

A Proofs

Proof of Lemma 2.1.

To define *consistency* more formally, I solve the consumer's lifetime maximisation problem in (1). Substituting the technology constraint $\tilde{z}_t = \tilde{A}\tilde{x}_t$ into period utility yields

$$\max_{\{\mathbf{x}_t, y_t\}_{t=1}^T} \sum_{t=1}^T \beta^{t-1} (u(\tilde{A}\tilde{x}_t) + y_t) \quad \text{subject to} \quad \sum_{t=1}^T \rho'_t \mathbf{x}_t + \sum_{t=1}^T \beta^{t-1} y_t = W,$$

where \mathbf{x}_0 is treated as fixed. Using the lifetime budget constraint to substitute out the outside good, the consumer's problem can be written as the unconstrained maximisation

$$\max_{\{\mathbf{x}_t\}_{t=1}^T} \sum_{t=1}^T \beta^{t-1} u(\tilde{A}\tilde{x}_t) + W - \sum_{t=1}^T \rho'_t \mathbf{x}_t.$$

Since W enters only as an additive constant, it does not affect the maximising choice of $\{\mathbf{x}_t\}_{t=1}^T$. I therefore equivalently represent the problem in constrained form by reintroducing the lifetime budget constraint and an associated Lagrange multiplier. The consumer's problem thus reduces to

$$\max_{\{\mathbf{x}_t\}_{t=1}^T} \sum_{t=1}^T \beta^{t-1} u(\tilde{A}\tilde{x}_t) \quad \text{subject to} \quad \sum_{t=1}^T \rho'_t \mathbf{x}_t = W,$$

where W is interpreted as lifetime wealth net of outside-good consumption.

The associated Lagrangian is

$$\mathcal{L}(\{\mathbf{x}_t\}) = \sum_{t=1}^T \beta^{t-1} u(\tilde{A}\tilde{x}_t) - \left\{ \sum_{t=1}^T \rho'_t \mathbf{x}_t - W \right\}, \quad (9)$$

where I normalise $\lambda = 1$ without loss of generality.

To define the first-order necessary conditions for an interior solution to this constrained optimisation problem, I require several vector derivatives. Applying the chain rule for both scalar and vector functions (Felippa, 2004) and using the "denominator layout" as my notational choice I have,

$$\underbrace{\frac{\partial u(\tilde{A}\tilde{x}_t)}{\partial \mathbf{x}_t}}_{(K \times 1)} = \underbrace{\frac{\partial \tilde{x}_t}{\partial \mathbf{x}_t}}_{(K \times 2K)} \underbrace{\frac{\partial u(\tilde{z}_t)}{\partial \tilde{x}_t}}_{(2K \times 1)} = \underbrace{\frac{\partial \tilde{x}_t}{\partial \mathbf{x}_t}}_{(K \times 2K)} \underbrace{\frac{\partial \tilde{z}_t}{\partial \tilde{x}_t}}_{(2K \times (J+J_2))} \underbrace{\frac{\partial u(\tilde{z}_t)}{\partial \tilde{z}_t}}_{((J+J_2) \times 1)} \quad (10)$$

where, recalling my notation defined in (2) I have,

$$\frac{\partial \tilde{\mathbf{x}}_t}{\partial \mathbf{x}_t} = \left[\mathbf{I}_{K \times K} \mid \mathbf{0}_{K \times K} \right],$$

$$\frac{\partial \tilde{\mathbf{z}}_t}{\partial \tilde{\mathbf{x}}_t} = \frac{\partial(\tilde{\mathbf{A}}\tilde{\mathbf{x}}_t)}{\partial \tilde{\mathbf{x}}_t} = \tilde{\mathbf{A}}',$$

$$\frac{\partial u(\tilde{\mathbf{z}}_t)}{\partial \tilde{\mathbf{z}}_t} = \partial u(\tilde{\mathbf{z}}_t) := \left[\partial_{\mathbf{z}_t^c} u(\tilde{\mathbf{z}}_t)', \partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_t)', \partial_{\mathbf{z}_{t-1}^a} u(\tilde{\mathbf{z}}_t)' \right]',$$

where \mid denotes the horizontal concatenation of the $K \times K$ identity matrix and the $K \times K$ matrix of zeros, and $\partial u(\tilde{\mathbf{z}})$ denotes the superderivative of u at $\tilde{\mathbf{z}}$. Repeating the chain rule exercise in (10), except this time differentiating with respect to the one period lag of market goods, I have,

$$\underbrace{\frac{\partial u(\tilde{\mathbf{A}}\tilde{\mathbf{x}}_t)}{\partial \mathbf{x}_{t-1}}}_{(K \times 1)} = \underbrace{\frac{\partial \tilde{\mathbf{x}}_t}{\partial \mathbf{x}_{t-1}}}_{(K \times 2K)} \underbrace{\frac{\partial u(\tilde{\mathbf{z}}_t)}{\partial \tilde{\mathbf{x}}_t}}_{(2K \times 1)} = \underbrace{\frac{\partial \tilde{\mathbf{x}}_t}{\partial \mathbf{x}_{t-1}}}_{(K \times 2K)} \underbrace{\frac{\partial \tilde{\mathbf{z}}_t}{\partial \tilde{\mathbf{x}}_t}}_{(2K \times (J+J_2))} \underbrace{\frac{\partial u(\tilde{\mathbf{z}}_t)}{\partial \tilde{\mathbf{z}}_t}}_{((J+J_2) \times 1)} \quad (11)$$

where the only new term is,

$$\frac{\partial \tilde{\mathbf{x}}_t}{\partial \mathbf{x}_{t-1}} = \left[\mathbf{0}_{K \times K} \mid \mathbf{I}_{K \times K} \right].$$

It follows from these intermediate calculations of the vector derivatives that,¹³

$$\frac{\partial u(\tilde{\mathbf{A}}\tilde{\mathbf{x}}_t)}{\partial \mathbf{x}_t} = \left[\mathbf{I}_{K \times K} \mid \mathbf{0}_{K \times K} \right] \tilde{\mathbf{A}}' \partial u(\tilde{\mathbf{z}}_t)$$

and,

$$\frac{\partial u(\tilde{\mathbf{A}}\tilde{\mathbf{x}}_{t+1})}{\partial \mathbf{x}_t} = \left[\mathbf{0}_{K \times K} \mid \mathbf{I}_{K \times K} \right] \tilde{\mathbf{A}}' \partial u(\tilde{\mathbf{z}}_{t+1}).$$

For interior dates $t \in \{1, \dots, T-1\}$, the relevant first-order necessary conditions associated with the Lagrangian in (9) now follow immediately as,

$$\partial_{\mathbf{x}_t} \mathcal{L} = 0 \quad \Rightarrow \quad \boldsymbol{\rho}_t = \beta^{t-1} \left(\left[\mathbf{I}_{K \times K} \mid \mathbf{0}_{K \times K} \right] \tilde{\mathbf{A}}' \partial u(\tilde{\mathbf{z}}_t) + \beta \left[\mathbf{0}_{K \times K} \mid \mathbf{I}_{K \times K} \right] \tilde{\mathbf{A}}' \partial u(\tilde{\mathbf{z}}_{t+1}) \right). \quad (12)$$

But recall from (2) that $\tilde{\mathbf{A}}$ is a $2K \times (J + J_2)$ block matrix. Hence, these first-order conditions can be substantially simplified. Indeed, since I have conformable partitions of the block matrices,

$$\begin{aligned} \left[\mathbf{I}_{K \times K} \mid \mathbf{0}_{K \times K} \right] \tilde{\mathbf{A}}' &= \left[\mathbf{I}_{K \times K} \mid \mathbf{0}_{K \times K} \right] \begin{bmatrix} \mathbf{A}' & \mathbf{0}_{K \times J_2} \\ \mathbf{0}_{K \times J} & (\mathbf{A}^a)' \end{bmatrix} \\ &= \left[\mathbf{I}_{K \times K} \mathbf{A}' + \mathbf{0}_{K \times K} \mathbf{0}_{K \times J} \mid \mathbf{I}_{K \times K} \mathbf{0}_{K \times J_2} + \mathbf{0}_{K \times K} (\mathbf{A}^a)' \right] \\ &= \left[\mathbf{A}' \mid \mathbf{0}_{K \times J_2} \right]. \end{aligned}$$

¹³Since $u(\cdot)$ is assumed to be concave but not necessarily differentiable, each $\partial u(\tilde{\mathbf{z}}_t)$ denotes a supergradient. As such, expressions like those below should formally be interpreted as *set inclusions* rather than strict equalities—e.g., $0 \in \partial_{\mathbf{x}_t} \mathcal{L}$. When u is differentiable at $\tilde{\mathbf{z}}_t$, the supergradient is a singleton and the equality holds exactly.

Analogously,

$$\begin{aligned}
\beta \left[\mathbf{0}_{K \times K} \mid \mathbf{I}_{K \times K} \right] \tilde{\mathbf{A}}' &= \beta \left[\mathbf{0}_{K \times K} \mid \mathbf{I}_{K \times K} \right] \begin{bmatrix} \mathbf{A}' & \mathbf{0}_{K \times J_2} \\ \mathbf{0}_{K \times J} & (\mathbf{A}^a)^\prime \end{bmatrix} \\
&= \beta \left[\mathbf{0}_{K \times K} \mathbf{A}' + \mathbf{I}_{K \times K} \mathbf{0}_{K \times J} \mid \mathbf{0}_{K \times K} \mathbf{0}_{K \times J_2} + \mathbf{I}_{K \times K} (\mathbf{A}^a)^\prime \right] \\
&= \beta \left[\mathbf{0}_{K \times J} \mid (\mathbf{A}^a)^\prime \right].
\end{aligned}$$

Inserting these simplifications, the first-order conditions in (12) reduce to:

$$\boldsymbol{\rho}_t = \beta^{t-1} \left(\left[\mathbf{A}' \mid \mathbf{0}_{K \times J_2} \right] \partial u(\tilde{\mathbf{z}}_t) + \beta \left[\mathbf{0}_{K \times J} \mid (\mathbf{A}^a)^\prime \right] \partial u(\tilde{\mathbf{z}}_{t+1}) \right). \quad (13)$$

But again, since the supergradient $\partial u(\tilde{\mathbf{z}}_t)$ can be partitioned as a $J + J_2$ block vector,

$$\partial u(\tilde{\mathbf{z}}_t) := \begin{bmatrix} \left[\begin{array}{c} \partial_{\mathbf{z}_t^c} u(\tilde{\mathbf{z}}_t) \\ \partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_t) \end{array} \right] \\ \left[\partial_{\mathbf{z}_{t-1}^a} u(\tilde{\mathbf{z}}_t) \right] \end{bmatrix},$$

the first-order conditions in (13) further simplify to,

$$\begin{aligned}
\boldsymbol{\rho}_t &= \beta^{t-1} \left(\left[\mathbf{A}' \mid \mathbf{0}_{K \times J_2} \right] \begin{bmatrix} \left[\begin{array}{c} \partial_{\mathbf{z}_t^c} u(\tilde{\mathbf{z}}_t) \\ \partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_t) \end{array} \right] \\ \left[\partial_{\mathbf{z}_{t-1}^a} u(\tilde{\mathbf{z}}_t) \right] \end{bmatrix} + \beta \left[\mathbf{0}_{K \times J} \mid (\mathbf{A}^a)^\prime \right] \begin{bmatrix} \left[\begin{array}{c} \partial_{\mathbf{z}_{t+1}^c} u(\tilde{\mathbf{z}}_{t+1}) \\ \partial_{\mathbf{z}_{t+1}^a} u(\tilde{\mathbf{z}}_{t+1}) \end{array} \right] \\ \left[\partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_{t+1}) \right] \end{bmatrix} \right) \\
&= \beta^{t-1} \left(\mathbf{A}' \begin{bmatrix} \partial_{\mathbf{z}_t^c} u(\tilde{\mathbf{z}}_t) \\ \partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_t) \end{bmatrix} + \beta (\mathbf{A}^a)^\prime \left[\partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_{t+1}) \right] \right). \quad (14)
\end{aligned}$$

At the terminal date, the continuation term vanishes. Hence the corresponding first-order condition reduces to

$$\boldsymbol{\rho}_T = \beta^{T-1} \mathbf{A}' \begin{bmatrix} \partial_{\mathbf{z}_T^c} u(\tilde{\mathbf{z}}_T) \\ \partial_{\mathbf{z}_T^a} u(\tilde{\mathbf{z}}_T) \end{bmatrix}, \quad (15)$$

which is exactly the same pricing condition as in (14) under the convention $\boldsymbol{\pi}_{T+1}^1 \equiv \mathbf{0}$. Since I assume u to be concave, the associated KKT conditions are sufficient for a global maximum. Hence, allowing for corner solutions, I replace the stationarity equalities in (14) and (15) with inequalities to obtain the full set of price and data pairs $\{\boldsymbol{\rho}_t, \mathbf{x}_t\}_{t=1}^T$ consistent with an interior or corner solution to the consumer's maximisation problem. This gives rise to my formal definition of *consistency* in Definition 2.1 and Lemma 2.1. \square

Proof of Theorem 2.1.

(A) \Rightarrow (B): Assume (A) holds. Then, by Lemma 2.1 of consistency I have that for all $t \in \{1, \dots, T\}$,

$$\rho_t \geq \mathbf{A}' \boldsymbol{\pi}_t^0 + (\mathbf{A}^a)' \boldsymbol{\pi}_{t+1}^1,$$

with equality for all k such that $x_t^k > 0$, where $\boldsymbol{\pi}_{T+1}^1 \equiv \mathbf{0}$. Element wise, this states that for all k and for all $t \in \{1, \dots, T\}$,

$$\rho_t^k \geq \mathbf{a}'_k \boldsymbol{\pi}_t^0 + \mathbf{a}^{a'}_k \boldsymbol{\pi}_{t+1}^1,$$

and where $x_t^k > 0$,

$$\rho_t^k = \mathbf{a}'_k \boldsymbol{\pi}_t^0 + \mathbf{a}^{a'}_k \boldsymbol{\pi}_{t+1}^1.$$

This gives us restrictions (B2) and (B3), respectively.

It remains to show (B1) holds. Using the augmented notation $\tilde{\mathbf{z}}_t := ((\mathbf{z}_t^c)', (\mathbf{z}_t^a)', (\mathbf{z}_{t-1}^a)')'$ and $\tilde{\boldsymbol{\pi}}_t := \frac{1}{\beta^{t-1}} [\boldsymbol{\pi}_t^{0'}, \boldsymbol{\pi}_t^{1'}]'$ it follows from Definitions (3) and (4) for the shadow prices that for any element of the superdifferential of u at $\tilde{\mathbf{z}}_t$,

$$\begin{aligned} \partial u(\tilde{\mathbf{z}}_t)' &= \left[[\partial_{\mathbf{z}_t^c} u(\tilde{\mathbf{z}}_t)', \partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_t)'], \partial_{\mathbf{z}_{t-1}^a} u(\tilde{\mathbf{z}}_t)' \right] \\ &= \frac{1}{\beta^{t-1}} \left[\beta^{t-1} [\partial_{\mathbf{z}_t^c} u(\tilde{\mathbf{z}}_t)', \partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_t)'], \beta^{t-1} \partial_{\mathbf{z}_{t-1}^a} u(\tilde{\mathbf{z}}_t)' \right] \\ &= \frac{1}{\beta^{t-1}} [\boldsymbol{\pi}_t^{0'}, \boldsymbol{\pi}_t^{1'}] \\ &= \tilde{\boldsymbol{\pi}}_t'. \end{aligned}$$

The concavity and superdifferentiability of the instantaneous utility function $u(\tilde{\mathbf{z}}_t)$ means,

$$u(\tilde{\mathbf{z}}_s) - u(\tilde{\mathbf{z}}_t) \leq \partial u(\tilde{\mathbf{z}}_t)'(\tilde{\mathbf{z}}_s - \tilde{\mathbf{z}}_t) \quad \forall s, t \in \{1, \dots, T\}.$$

Combining this concavity result with that implied by the optimising behaviour above therefore implies,

$$u(\tilde{\mathbf{z}}_s) - u(\tilde{\mathbf{z}}_t) \leq \tilde{\boldsymbol{\pi}}_t'(\tilde{\mathbf{z}}_s - \tilde{\mathbf{z}}_t) \quad \forall s, t \in \{1, \dots, T\}.$$

Take any finite ordered cycle (t_1, \dots, t_M) from τ , with $t_{M+1} = t_1$. Summing the corresponding inequalities

$$u(\tilde{\mathbf{z}}_{t_{m+1}}) - u(\tilde{\mathbf{z}}_{t_m}) \leq \tilde{\boldsymbol{\pi}}_{t_m}'(\tilde{\mathbf{z}}_{t_{m+1}} - \tilde{\mathbf{z}}_{t_m}) \quad m = 1, \dots, M$$

gives

$$0 \leq \sum_{m=1}^M \tilde{\boldsymbol{\pi}}_{t_m}'(\tilde{\mathbf{z}}_{t_{m+1}} - \tilde{\mathbf{z}}_{t_m}),$$

which is restriction (B1).

(B) \Rightarrow (A): Restriction (B1) imposes that the finite dataset $(\tilde{\boldsymbol{\pi}}_t, \tilde{\mathbf{z}}_t), t = 1, \dots, T$ is cyclically monotone as defined by Browning (1989). By the standard finite-sample Afriat inequalities (Afriat, 1967; Diewert, 1973; Varian, 1982), this is equivalent to the existence of T numbers $\{V_t\}_{t=1}^T$ such that

$$V_s \leq V_t + \tilde{\boldsymbol{\pi}}_t'(\tilde{\mathbf{z}}_s - \tilde{\mathbf{z}}_t) \quad \forall s, t \in \tau.$$

Applying Afriat's theorem to the finite dataset $\{\tilde{\boldsymbol{\pi}}_t, \tilde{\mathbf{z}}_t\}_{t \in \tau}$ then implies that there exists a locally non-

satiated and concave utility function rationalising these observations. One such representation is the Afriat envelope,

$$u(\tilde{\mathbf{z}}) := \min_{t \in \tau} \{V_t + \tilde{\boldsymbol{\pi}}_t'(\tilde{\mathbf{z}} - \tilde{\mathbf{z}}_t)\}.$$

Since $u(\cdot)$ is the minimum of finitely many affine functions, it is proper, concave, and superdifferentiable. Moreover, the Afriat inequalities imply that at each observed $\tilde{\mathbf{z}}_t$ the t -th affine function supports u , so that

$$\tilde{\boldsymbol{\pi}}_t \in \partial u(\tilde{\mathbf{z}}_t) \quad \forall t \in \tau.$$

Selecting one such supergradient at each observed point and partitioning it conformably yields the following relationships between shadow prices and supergradients of the utility function:

$$\boldsymbol{\pi}_t^0 = \beta^{t-1} \begin{bmatrix} \partial_{\mathbf{z}_t^c} u(\tilde{\mathbf{z}}_t) \\ \partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_t) \end{bmatrix} \quad (16)$$

$$\boldsymbol{\pi}_t^1 = \beta^{t-1} \left[\partial_{\mathbf{z}_{t-1}^a} u(\tilde{\mathbf{z}}_t) \right]. \quad (17)$$

Since (17) holds for all $t \in \tau$, this condition can be forwarded one period to obtain,

$$\boldsymbol{\pi}_{t+1}^1 = \beta^t \left[\partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_{t+1}) \right] \quad (18)$$

for all $t \in \{1, \dots, T-1\}$. Substituting expressions (16) and (18) into condition (B2) re-written in matrix form, $\boldsymbol{\rho}_t \geq \mathbf{A}'\boldsymbol{\pi}_t^0 + (\mathbf{A}^a)'\boldsymbol{\pi}_{t+1}^1$, I obtain,

$$\boldsymbol{\rho}_t \geq \mathbf{A}'\beta^{t-1} \begin{bmatrix} \partial_{\mathbf{z}_t^c} u(\tilde{\mathbf{z}}_t) \\ \partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_t) \end{bmatrix} + (\mathbf{A}^a)'\beta^t \left[\partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_{t+1}) \right]. \quad (19)$$

At the terminal date, condition (B2) reduces under $\boldsymbol{\pi}_{T+1}^1 \equiv \mathbf{0}$ to

$$\boldsymbol{\rho}_T \geq \mathbf{A}'\beta^{T-1} \begin{bmatrix} \partial_{\mathbf{z}_T^c} u(\tilde{\mathbf{z}}_T) \\ \partial_{\mathbf{z}_T^a} u(\tilde{\mathbf{z}}_T) \end{bmatrix}.$$

In the event that a good k is consumed in strictly positive amounts, then substituting expressions (16) and (18) into condition (B3) yields the same expression as in (19), but with equality; the terminal-date analogue again follows by setting $\boldsymbol{\pi}_{T+1}^1 \equiv \mathbf{0}$. By Lemma 2.1, this means that the data $\{\boldsymbol{\rho}_t; \mathbf{x}_t\}_{t \in \{1, \dots, T\}}$ are consistent with the one-lag habits model for given technology \mathbf{A} . \square

Proof of Theorem 2.2.

$(A^+) \Rightarrow (B^+)$: Given $\{\boldsymbol{\rho}_t^0\}_{t \in \{1, \dots, T\}}$ I have a full set of prices, $\{\boldsymbol{\rho}_t\}_{t \in \{1, \dots, T\}}$, such that the data satisfies the model. Hence, by Theorem 2.1, condition (B) holds. Hence, (B^+) also holds.

$(B^+) \Rightarrow (A^+)$: I have shadow discounted prices $\{\boldsymbol{\pi}_t^0\}_{t \in \{1, \dots, T\}}$ and $\{\boldsymbol{\pi}_t^1\}_{t \in \{1, \dots, T\}}$ such that $(B1^+)$ and $(B2^+)$ hold. Use these shadow prices to construct the unobserved prices via,

$$\boldsymbol{\rho}_t^0 = (\mathbf{B}_t^0)'\boldsymbol{\pi}_t^0 + (\mathbf{B}_t^{a,0})'\boldsymbol{\pi}_{t+1}^1.$$

with $\pi_{T+1}^1 \equiv \mathbf{0}$. Using these constructed prices, it follows by comparison with Theorem 2.1 that (A⁺) holds. \square

Habits over goods as a special case of the characteristics model (Definition 3.1).

The (one-lag) habits-over-characteristics model nests the habits-over-goods model of Crawford (2010) as a special case. To see this, consider a trivial characteristics model where $J = K$ and the technology matrix \mathbf{A} is the $J \times J$ identity matrix. Then:

$$\mathbf{z}_t = \begin{bmatrix} \mathbf{z}_t^c \\ \mathbf{z}_t^a \end{bmatrix} = \begin{bmatrix} \mathbf{x}_t^c \\ \mathbf{x}_t^a \end{bmatrix} = \mathbf{x}_t,$$

where \mathbf{x}_t^c and \mathbf{x}_t^a denote the J_1 non-habit-forming and J_2 habit-forming goods, respectively. Under this assumption, the matrices $(\mathbf{A}^c)'$ and $(\mathbf{A}^a)'$ take on block-structured forms. The matrix $(\mathbf{A}^c)'$ is a $K \times J_1$ matrix whose first J_1 rows form a $J_1 \times J_1$ identity matrix, with all remaining entries equal to zero. Conversely, $(\mathbf{A}^a)'$ is a $K \times J_2$ matrix whose bottom J_2 rows comprise a $J_2 \times J_2$ identity matrix, while the entries in the first $K - J_2$ rows are zero.

Substituting the identity structure of \mathbf{A} into Equation (\star) and defining the augmented bundle $\bar{\mathbf{x}}_t := (\mathbf{x}_t^c, \mathbf{x}_t^a, \mathbf{x}_{t-1}^a)'$ = $\tilde{\mathbf{z}}_t$, the first-order condition becomes,

$$\boldsymbol{\rho}_t = \begin{bmatrix} \boldsymbol{\rho}_t^c \\ \boldsymbol{\rho}_t^a \end{bmatrix} \geq \begin{bmatrix} \boldsymbol{\pi}_t^{c,0} \\ \mathbf{0}_{K-J_1} \end{bmatrix} + \begin{bmatrix} \mathbf{0}_{K-J_2} \\ \boldsymbol{\pi}_t^{a,0} \end{bmatrix} + \begin{bmatrix} \mathbf{0}_{K-J_2} \\ \boldsymbol{\pi}_{t+1}^1 \end{bmatrix} = \beta^{t-1} \begin{bmatrix} \partial_{\mathbf{x}_t^c} u(\bar{\mathbf{x}}_t) \\ \mathbf{0}_{K-J_1} \end{bmatrix} + \beta^{t-1} \begin{bmatrix} \mathbf{0}_{K-J_2} \\ \partial_{\mathbf{x}_t^a} u(\bar{\mathbf{x}}_t) \end{bmatrix} + \beta^t \begin{bmatrix} \mathbf{0}_{K-J_2} \\ \partial_{\mathbf{x}_t^a} u(\bar{\mathbf{x}}_{t+1}) \end{bmatrix},$$

where $\boldsymbol{\pi}_t^{c,0}$, $\boldsymbol{\pi}_t^{a,0}$ and $\boldsymbol{\rho}_t^c$, $\boldsymbol{\rho}_t^a$ are the subvectors of $\boldsymbol{\pi}_t^0$ and $\boldsymbol{\rho}_t$ corresponding to non-habit-forming and habit-forming goods, respectively, and the final equality follows using Definitions 3 and 4. For all k such that $x_t^k > 0$, these inequalities hold with equality. Thus, under the restriction $J = K$ and $\mathbf{A} = \mathbf{I}_J$, Definition 2.1 simplifies as in Definition 3.1.

Proof of Corollary 3.1.

Follows directly from Theorem 2.1 under the restriction $J = K$ and $\mathbf{A} = \mathbf{I}_J$. Equivalently, one may derive it from Definition 3.1 by replicating the proof of Theorem 2.1. \square

Intertemporally separable preferences over characteristics (Definition 3.2).

To formalise this nesting result, assume $\mathbf{z}_t = \mathbf{z}_t^c$ and embed the model in a lifecycle framework with a single intertemporal budget constraint and maintain the normalisation that the marginal utility of lifetime wealth equals one. Then the shadow price in Equation (3) reduces to:

$$\boldsymbol{\pi}_t = \partial_{\mathbf{z}_t} u(\mathbf{z}_t) \tag{20}$$

noting that β no longer appears due to intertemporal separability.¹⁴ Under these assumptions, Definition 2.1 reduces to that in Definition 3.2.

Proof of Corollary 3.2.

Corollary 3.2 is immediate from Theorem 2.1 upon setting $J_1 = J$ and $\beta = 1$.

¹⁴This is equivalent to setting $\beta = 1$, which is without loss of generality in this setting.

I now show that this Corollary is equivalent to Theorem 1 of Blow et al. (2008) when the static characteristics model is embedded in a lifecycle framework. Equivalence between (A'') and the P'' condition in their theorem follows directly from Definition 3.2. Equivalence of conditions (B2'') and (B3'') with (A2) and (A3) in Blow et al. (2008) is straightforward after aligning notation.

The key step is to show that condition (B1'') is equivalent to condition (A1) in Blow et al. (2008), which states that there exist T scalars V_t and T vectors $\boldsymbol{\pi}_t$ such that:

$$V_s \leq V_t + \lambda_t \boldsymbol{\pi}'_t (\mathbf{A}\mathbf{x}_s - \mathbf{A}\mathbf{x}_t) \quad \forall s, t. \quad (\text{A1})$$

In Blow et al. (2008), λ_t denotes the (time- t) marginal utility of wealth. Under my single lifetime budget constraint I have $\lambda_t \equiv \lambda$, and I maintain the normalisation $\lambda = 1$ throughout.

To show that (B1'') implies (A1), note that (B1'') imposes cyclical monotonicity on the finite dataset $\{\boldsymbol{\pi}_t, \mathbf{z}_t\}_{t=1}^T$. By the standard finite-sample Afriat inequalities (Afriat, 1967; Diewert, 1973; Varian, 1982), this is equivalent to the existence of T numbers $\{V_t\}_{t=1}^T$ such that

$$V_s \leq V_t + \boldsymbol{\pi}'_t (\mathbf{z}_s - \mathbf{z}_t) \quad \forall s, t.$$

Since $\mathbf{z}_t = \mathbf{A}\mathbf{x}_t$ in the static characteristics model, this becomes

$$V_s \leq V_t + \boldsymbol{\pi}'_t (\mathbf{A}\mathbf{x}_s - \mathbf{A}\mathbf{x}_t) \quad \forall s, t,$$

which is (A1) under the normalisation $\lambda_t \equiv 1$.

Conversely, suppose (A1) holds. Under $\lambda = 1$, rearranging yields:

$$0 \leq (V_t - V_s) + \boldsymbol{\pi}'_t (\mathbf{A}\mathbf{x}_s - \mathbf{A}\mathbf{x}_t).$$

Take any finite ordered cycle (t_1, \dots, t_M) with $t_{M+1} = t_1$. Summing the corresponding inequalities over $m = 1, \dots, M$ makes the Afriat numbers telescope, yielding

$$0 \leq \sum_{m=1}^M \boldsymbol{\pi}'_{t_m} (\mathbf{A}\mathbf{x}_{t_{m+1}} - \mathbf{A}\mathbf{x}_{t_m}),$$

which is (B1'').

I conclude that Corollary 3.2 provides an equivalent characterisation to Blow et al. (2008) when the static characteristics model is embedded in a lifecycle framework. This highlights the additional empirical restrictions imposed by intertemporal optimisation: a consumer consistent with intratemporal utility maximisation over characteristics must also allocate expenditure across periods to maximise lifetime utility. \square

B Non-linear characteristics model

B.1 Consumer problem in the non-linear model

This paper focuses on the linear characteristics model, $\mathbf{z}_t = \mathbf{A}\mathbf{x}_t$. However, most of my analysis generalises to a non-linear characteristics setting where I assume $\mathbf{z}_t = \mathbf{F}(\mathbf{x}_t)$ for some concave, increasing function $\mathbf{F} : \mathbb{R}_{\geq 0}^K \rightarrow \mathbb{R}_{\geq 0}^J$. Here, I provide an analogue notion of consistency and the relevant Afriat Theorem for such non-linear technologies. For convenience, I take \mathbf{F} to be differentiable, with an associated (denominator layout) $K \times J$ matrix derivative at \mathbf{x}_t given by,

$$\nabla \mathbf{F}(\mathbf{x}_t) := \left(\frac{\partial \mathbf{z}_t^c}{\partial \mathbf{x}_t} \quad \bigg| \quad \frac{\partial \mathbf{z}_t^a}{\partial \mathbf{x}_t} \right) = \begin{pmatrix} \frac{\partial z_{1,t}^c}{\partial x_{1,t}} & \cdots & \frac{\partial z_{J_1,t}^c}{\partial x_{1,t}} & \frac{\partial z_{1,t}^a}{\partial x_{1,t}} & \cdots & \frac{\partial z_{J_2,t}^a}{\partial x_{1,t}} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial z_{1,t}^c}{\partial x_{K,t}} & \cdots & \frac{\partial z_{J_1,t}^c}{\partial x_{K,t}} & \frac{\partial z_{1,t}^a}{\partial x_{K,t}} & \cdots & \frac{\partial z_{J_2,t}^a}{\partial x_{K,t}} \end{pmatrix}.$$

If instead \mathbf{F} is only superdifferentiable, one must replace all gradients with the superdifferential.

Use the augmented notation for $\tilde{\mathbf{x}}_t$ and $\tilde{\mathbf{z}}_t$ from Section 2. In addition, define the augmented technology function $\tilde{\mathbf{F}} : \mathbb{R}_{\geq 0}^{2K} \rightarrow \mathbb{R}_{\geq 0}^{J+J_2}$ via,

$$\tilde{\mathbf{F}}(\tilde{\mathbf{x}}_t) := \begin{pmatrix} \mathbf{F}(\mathbf{x}_t) \\ \mathbf{F}^a(\mathbf{x}_{t-1}) \end{pmatrix}, \quad (21)$$

where $\mathbf{F}^a : \mathbb{R}_{\geq 0}^K \rightarrow \mathbb{R}_{\geq 0}^{J_2}$ is the function defined by taking the last J_2 components of \mathbf{F} . Note that the derivative of this extended vector-valued function is a $2K \times (J + J_2)$ matrix given by,

$$\nabla \tilde{\mathbf{F}}(\tilde{\mathbf{x}}_t) = \begin{pmatrix} \frac{\partial \mathbf{z}_t^c}{\partial \mathbf{x}_t} & \bigg| & \frac{\partial \mathbf{z}_t^a}{\partial \mathbf{x}_t} & \bigg| & \mathbf{0}_{K \times J_2} \\ \mathbf{0}_{K \times J_1} & \bigg| & \mathbf{0}_{K \times J_2} & \bigg| & \frac{\partial \mathbf{z}_{t-1}^a}{\partial \mathbf{x}_{t-1}} \end{pmatrix} = \begin{pmatrix} \nabla \mathbf{F}(\mathbf{x}_t) & \bigg| & \mathbf{0}_{K \times J_2} \\ \mathbf{0}_{K \times J} & \bigg| & \nabla \mathbf{F}^a(\mathbf{x}_{t-1}) \end{pmatrix} \quad (22)$$

where $\nabla \mathbf{F}^a(\mathbf{x}_{t-1})$ is the $K \times J_2$ submatrix found by taking the last J_2 columns of $\nabla \mathbf{F}(\mathbf{x}_{t-1})$ and I use the fact that $\frac{\partial \mathbf{z}_{t-1}^a}{\partial \mathbf{x}_t} = \mathbf{0}_{K \times J_2}$ and $\left(\frac{\partial \mathbf{z}_t^c}{\partial \mathbf{x}_{t-1}} \quad \bigg| \quad \frac{\partial \mathbf{z}_t^a}{\partial \mathbf{x}_{t-1}} \right) = \left(\mathbf{0}_{K \times J_1} \quad \bigg| \quad \mathbf{0}_{K \times J_2} \right)$.

Using this augmented notation, my quasilinear model of interest becomes:

$$\max_{\{\mathbf{x}_t, y_t\}_{t=1}^T} \sum_{t=1}^T \beta^{t-1} (u(\tilde{\mathbf{z}}_t) + y_t) \quad \text{subject to} \quad \sum_{t=1}^T \rho'_t \mathbf{x}_t + \sum_{t=1}^T \beta^{t-1} y_t = W, \quad \tilde{\mathbf{z}}_t = \tilde{\mathbf{F}}(\tilde{\mathbf{x}}_t). \quad (23)$$

By quasi-linearity, the outside good can be suppressed and the analysis can be conducted in terms of $\{\mathbf{x}_t\}$ and the present-value expenditure constraint; see Appendix A for details on this suppression step.

B.2 Consistency in the non-linear model

I now formalise the notion of *consistency* when the transformation technology is non-linear. As in Section 2, this amounts to solving the consumer's constrained maximisation problem defined in (23). Indeed, combining the technology constraint $\tilde{\mathbf{z}}_t = \tilde{\mathbf{F}}(\tilde{\mathbf{x}}_t)$ with quasi-linearity allows the outside good to be suppressed (Appendix A). The consumer's problem can therefore be written as

$$\max_{\{\mathbf{x}_t\}_{t=1}^T} \sum_{t=1}^T \beta^{t-1} u(\tilde{\mathbf{F}}(\tilde{\mathbf{x}}_t)) \quad \text{subject to} \quad \sum_{t=1}^T \rho'_t \mathbf{x}_t = W,$$

where W is now interpreted as lifetime wealth net of outside-good consumption. The associated Lagrangian is

$$\mathcal{L}(\{\mathbf{x}_t\}) = \sum_{t=1}^T \beta^{t-1} u(\tilde{\mathbf{F}}(\tilde{\mathbf{x}}_t)) - \left\{ \sum_{t=1}^T \boldsymbol{\rho}'_t \mathbf{x}_t - W \right\}, \quad (24)$$

where I normalise $\lambda = 1$ without loss of generality.

The first-order necessary conditions then follow analogously to Section 2, except that the derivative of the augmented characteristic vector with respect to the augmented market goods vector is now,

$$\frac{\partial \tilde{\mathbf{z}}_t}{\partial \tilde{\mathbf{x}}_t} = \frac{\partial \tilde{\mathbf{F}}(\tilde{\mathbf{x}}_t)}{\partial \tilde{\mathbf{x}}_t} = \nabla \tilde{\mathbf{F}}(\tilde{\mathbf{x}}_t),$$

as defined in (22). Hence, following the same simplification steps as in Section 2, the first-order conditions reduce to,

$$\boldsymbol{\rho}_t = \beta^{t-1} \left(\nabla \mathbf{F}(\mathbf{x}_t) \begin{bmatrix} \partial_{\mathbf{z}_t^c} u(\tilde{\mathbf{z}}_t) \\ \partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_t) \end{bmatrix} + \beta \nabla \mathbf{F}^a(\mathbf{x}_t) \begin{bmatrix} \partial_{\mathbf{z}_t^c} u(\tilde{\mathbf{z}}_{t+1}) \\ \partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_{t+1}) \end{bmatrix} \right), \quad (25)$$

where $\nabla \mathbf{F}^a(\mathbf{x}_t)$ is the $K \times J_2$ submatrix found by taking the last J_2 columns of $\nabla \mathbf{F}(\mathbf{x}_t)$. This gives rise to my formal definition of *consistency* in the non-linear model as follows.

Definition B.1. The data $\{\boldsymbol{\rho}_t; \mathbf{x}_t\}_{t \in \{1, \dots, T\}}$ are said to be *consistent* with the non-linear one-lag habits-over-characteristics model given the increasing, concave technology \mathbf{F} if they solve the agent's lifetime utility maximisation problem defined in Equation (23), for some locally non-satiated, increasing, superdifferentiable, and concave utility function $u(\cdot)$ and discount factor $\beta \in (0, 1]$.

The following lemma provides a set of necessary and sufficient conditions for this non-linear form of consistency to hold.

Lemma B.1. The data $\{\boldsymbol{\rho}_t; \mathbf{x}_t\}_{t \in \{1, \dots, T\}}$ are *consistent* with the non-linear one-lag habits-over-characteristics model given the increasing, concave technology \mathbf{F} if there exists a locally non-satiated, increasing, superdifferentiable, and concave utility function $u(\cdot)$ and a discount factor $\beta \in (0, 1]$ such that for all $t \in \{1, \dots, T-1\}$,

$$\boldsymbol{\rho}_t \geq \nabla \mathbf{F}(\mathbf{x}_t) \boldsymbol{\pi}_t^0 + \nabla \mathbf{F}^a(\mathbf{x}_t) \boldsymbol{\pi}_{t+1}^1, \quad (\star_N)$$

with equality for all k such that $x_t^k > 0$, and where discounted shadow prices are defined as:

$$\boldsymbol{\pi}_t^0 = \beta^{t-1} \begin{bmatrix} \partial_{\mathbf{z}_t^c} u(\tilde{\mathbf{z}}_t) \\ \partial_{\mathbf{z}_t^a} u(\tilde{\mathbf{z}}_t) \end{bmatrix}, \quad (SP_0)$$

$$\boldsymbol{\pi}_t^1 = \beta^{t-1} \begin{bmatrix} \partial_{\mathbf{z}_{t-1}^c} u(\tilde{\mathbf{z}}_t) \\ \partial_{\mathbf{z}_{t-1}^a} u(\tilde{\mathbf{z}}_t) \end{bmatrix}. \quad (SP_1)$$

where $\tilde{\mathbf{z}}_t = \tilde{\mathbf{F}}(\tilde{\mathbf{x}}_t)$ for all $t \in \{1, \dots, T\}$ and $\boldsymbol{\rho}_t$ denotes the vector of present-value prices. Since $u(\cdot)$ is increasing, the associated shadow prices satisfy $\boldsymbol{\pi}_t^0 \geq \mathbf{0}$ and $\boldsymbol{\pi}_t^1 \geq \mathbf{0}$ for all t .

Clearly, Definition B.1 and Lemma B.1 nest the linear one-lag habits-over-characteristics model when $\mathbf{F}(\mathbf{x}_t) = \mathbf{A}\mathbf{x}_t$. The key difference when $\mathbf{F}(\mathbf{x}_t) \neq \mathbf{A}\mathbf{x}_t$ is that the marginal product of the market goods in terms of the characteristics is no longer independent of demand. Prices therefore remain linear in the shadow discounted prices conditional on \mathbf{x}_t , but the coefficients in that linear combination now vary with demand through the Jacobian matrices. Using this more general, non-linear notion of consistency, I now derive testable empirical conditions involving only observables.

B.3 Afriat conditions in the non-linear model

I now generalise my main Theorem 2.1 to the non-linear model.

Theorem B.1. The following statements are equivalent:

- (A_N) The data $\{\boldsymbol{\rho}_t; \mathbf{x}_t\}_{t \in \{1, \dots, T\}}$ are consistent with the one-lag habits model given the increasing, concave technology \mathbf{F} .
- (B_N) There exist T J -vector shadow discounted prices $\{\boldsymbol{\pi}_t^0\}_{t \in \{1, \dots, T\}}$, T J_2 -vector shadow discounted prices $\{\boldsymbol{\pi}_t^1\}_{t \in \{1, \dots, T\}}$ and a discount factor $\beta \in (0, 1]$ such that,

$$0 \leq \sum_{\forall s, t \in \sigma} \tilde{\boldsymbol{\pi}}_s' (\tilde{\mathbf{z}}_t - \tilde{\mathbf{z}}_s) \quad \forall \sigma \subseteq \{1, \dots, T\} \quad (B1_N)$$

$$\rho_t^k \geq [\nabla \mathbf{F}(\mathbf{x}_t)]_k \boldsymbol{\pi}_t^0 + [\nabla \mathbf{F}^a(\mathbf{x}_t)]_k \boldsymbol{\pi}_{t+1}^1 \quad \forall k, t \in \{1, \dots, T-1\} \quad (B2_N)$$

$$\rho_t^k = [\nabla \mathbf{F}(\mathbf{x}_t)]_k \boldsymbol{\pi}_t^0 + [\nabla \mathbf{F}^a(\mathbf{x}_t)]_k \boldsymbol{\pi}_{t+1}^1 \quad \text{if } x_t^k > 0, \forall k, t \in \{1, \dots, T-1\} \quad (B3_N)$$

where $[\nabla \mathbf{F}(\mathbf{x}_t)]_k$ is the J -row vector corresponding to the k -th row of $\nabla \mathbf{F}(\mathbf{x}_t)$, $[\nabla \mathbf{F}^a(\mathbf{x}_t)]_k$ is the J_2 -row vector corresponding to the k -th row of $\nabla \mathbf{F}^a(\mathbf{x}_t)$, and $\tilde{\boldsymbol{\pi}}_t := \frac{1}{\beta^{t-1}} [\boldsymbol{\pi}_t^0, \boldsymbol{\pi}_t^1]'$, with $\boldsymbol{\pi}_t^0 \geq \mathbf{0}$ and $\boldsymbol{\pi}_t^1 \geq \mathbf{0}$ for all t .

Proof. The argument is identical to that for Theorem 2.1, using the updated notion of consistency given in Definition B.1. The only additional point is that, because \mathbf{F} is concave and increasing and u is concave and increasing, the composite objective $u(\tilde{\mathbf{F}}(\tilde{\mathbf{x}}_t))$ is concave in $\tilde{\mathbf{x}}_t$. Hence the first-order conditions in Lemma B.1 are sufficient for global optimality, and the remainder of the proof goes through unchanged. \square

C Additional data description and summary statistics

This appendix reports additional descriptive statistics for the scanner panel, including sample selection details, purchase intensity, brand concentration, price distributions, and characteristics distributions.

C.1 Sample selection

I begin by documenting the construction of the estimation sample from the raw 2010–2011 scanner panel. Table 7 reports the sequential filtering steps and resulting sample sizes. The first four restrictions impose minimum activity and panel-length requirements required for feasibility of the revealed preference framework and are therefore purely mechanical. The final restriction—dropping households whose purchased UPCs lack complete characteristics data—is the only potentially selective cut and is evaluated separately using balance tests.

Table 7: Sample construction and filtering

Step	Households	UPCs (K)	Purely mechanical
raw scanner (2010–11 combined)	4697	1139	Yes
drop HHs with only one observed purchase occasion	4505	1136	Yes
drop HHs with only one constructed time period ($T_i \leq 1$)	3673	1131	Yes
drop HHs with $T_i \leq 2$ (not meaningful for habit formation)	2972	1121	Yes
drop HHs who purchase UPCs with missing characteristics data	2282	801	No

Notes: Rows report sequential sample restrictions applied to the 2010–11 scanner panel. T_i denotes the number of household-specific constructed time periods. The restriction $T_i \leq 2$ ensures at least two transitions for identification of one-lag habit formation. “Purely mechanical” indicates whether the restriction is driven by panel structure rather than missing product characteristics.

As shown in Table 7, most sample attrition occurs due to mechanical feasibility restrictions, while the potentially selective exclusion due to missing characteristics data affects a comparatively smaller subset of households.

To assess whether exclusion due to missing characteristics data induces observable selection, Table 8 compares demographic characteristics between the baseline sample—defined as households satisfying all mechanical feasibility restrictions—and the final analysis sample. Reported p-values correspond to Pearson chi-squared tests of equality in distributions across groups.

Across most dimensions, observable characteristics are similar between excluded and retained households. While some differences arise along geographic location and household composition, these differences are modest in magnitude, and joint tests fail to reject equality across the majority of demographic characteristics. Overall, the balance results suggest limited scope for selection on observables arising from the final exclusion.

Table 8: Demographic balance: excluded versus final sample

Characteristic	Overall N = 2,972	Excluded N = 690	Final Sample N = 2,282	p-value
Family size				0.012
1	550 (19%)	118 (17%)	432 (19%)	
2	1,362 (46%)	293 (42%)	1,069 (47%)	
3+	1,060 (36%)	279 (40%)	781 (34%)	
Marital status				> 0.9
Married	307 (10%)	70 (10%)	237 (10%)	
Not married	2,658 (90%)	617 (90%)	2,041 (90%)	
Region				< 0.001
Eau Claire	1,597 (54%)	278 (40%)	1,319 (58%)	
Pittsfield	1,375 (46%)	412 (60%)	963 (42%)	
Children				0.032
Any children	667 (22%)	176 (26%)	491 (22%)	
No children	2,305 (78%)	514 (74%)	1,791 (78%)	
Occupation				0.9
Blue collar/service	1,149 (39%)	267 (39%)	882 (39%)	
Office/sales (sales/clerical)	454 (15%)	100 (14%)	354 (16%)	
Other/unknown	215 (7.2%)	55 (8.0%)	160 (7.0%)	
Retired	254 (8.5%)	62 (9.0%)	192 (8.4%)	
White collar (prof/manager)	900 (30%)	206 (30%)	694 (30%)	
Education				0.8
≤ High school	1,737 (61%)	400 (60%)	1,337 (61%)	
College+	203 (7.1%)	46 (6.9%)	157 (7.2%)	
Some college/tech	911 (32%)	219 (33%)	692 (32%)	

Notes: Counts are n (%). The “Excluded” column reports households removed due to missing UPC characteristic information; all other panel-length restrictions are satisfied in both groups. p-values are from Pearson chi-squared tests of equality across excluded and retained households.

C.2 Purchase intensity and brand concentration


Figure 4 shows the distribution of purchase counts per household-period. Figure 5 reports the distributions of average period consumption and expenditure. Figure 6 reports the distribution of the number of distinct brands purchased per period.

Figure 4: Units purchased per household-period



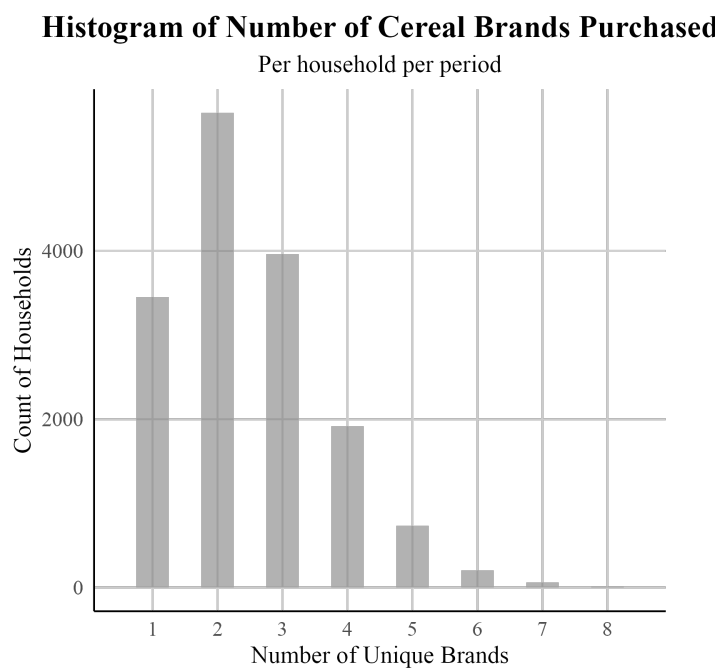
Notes: Histogram of total cereal units purchased by household i in period t , pooled across all households and constructed time periods.

Figure 5: Average purchasing intensity across households



Notes: Distributions of household-level averages of units purchased per period and expenditure per period across the sample.

Figure 6: Distinct brands purchased per household-period




Notes: Histogram of the number of distinct cereal brands purchased by a household within a constructed time period.

C.3 Price distributions

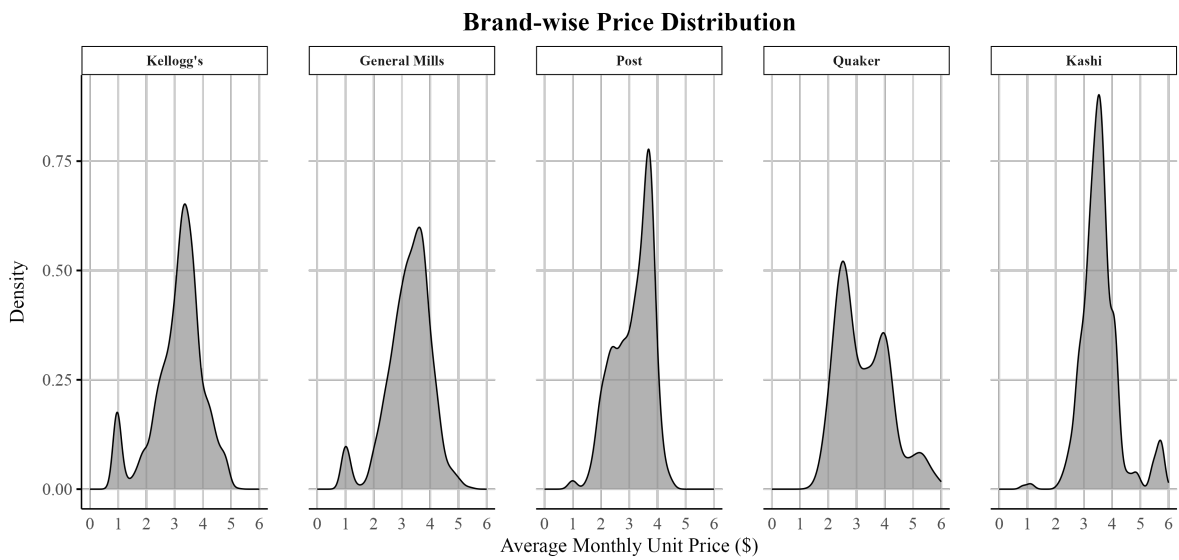
Figure 7 plots the distribution of average unit values across UPCs. Figure 8 shows the corresponding distributions for the five most prevalent brands.

Figure 7: Average unit values by UPC



Notes: Distribution of average unit values (expenditure divided by quantity) computed at the UPC level over the sample period.

Figure 8: Average unit values by brand



Notes: Distribution of UPC-level average unit values for the five most prevalent cereal brands in the sample.

C.4 Characteristics summary statistics

Tables 9 and 10 report descriptive statistics for the 23 characteristics used in the analysis. Nutritional attributes are expressed per 100g serving. Calories cluster tightly across products, while carbohydrates and sugar are generally high and fat content is low. Sodium exhibits substantial cross-product variation. Among binary attributes, whole-grain indicators and health-related descriptors are common, whereas organic and gluten-free labels are rare. Brand indicators reflect the dominance of Kellogg’s and General Mills in the sample.

Table 9: Nutritional characteristics of cereal products (per 100g)

Characteristic	units	mean	sd	q1	median	q3
calories	kcal	379	38	364	379	400
carbohydrates	g	81	7	78	82	85
total fat	g	4	4	2	3	5
saturated fat	g	1	1	0	0	1
fiber	g	8	5	3	7	10
protein	g	8	3	6	7	10
sodium	mg	435	239	277	467	600
sugar	g	25	12	18	25	33

Notes: Summary statistics computed across the UPCs in the final sample after excluding two products with missing serving-size entries. Nutritional values are standardised per 100g serving and rounded to the nearest integer.

Table 10: Binary product characteristics

Characteristic	Share (%)	Characteristic	Share (%)
whole grain	57	honey flavour	15
organic	8	gluten-free	3
oat-based	24	Kellogg’s	18
granola	10	General Mills	13
health halos	33	Post	11
fruity	6	Quaker	7
nutty	11	Kashi	4
chocolatey	5		

Notes: Shares denote the percentage of the $K = 801$ UPCs in the final sample exhibiting the indicated characteristic. Values are rounded to the nearest integer.

D Identification of the Discount Factor

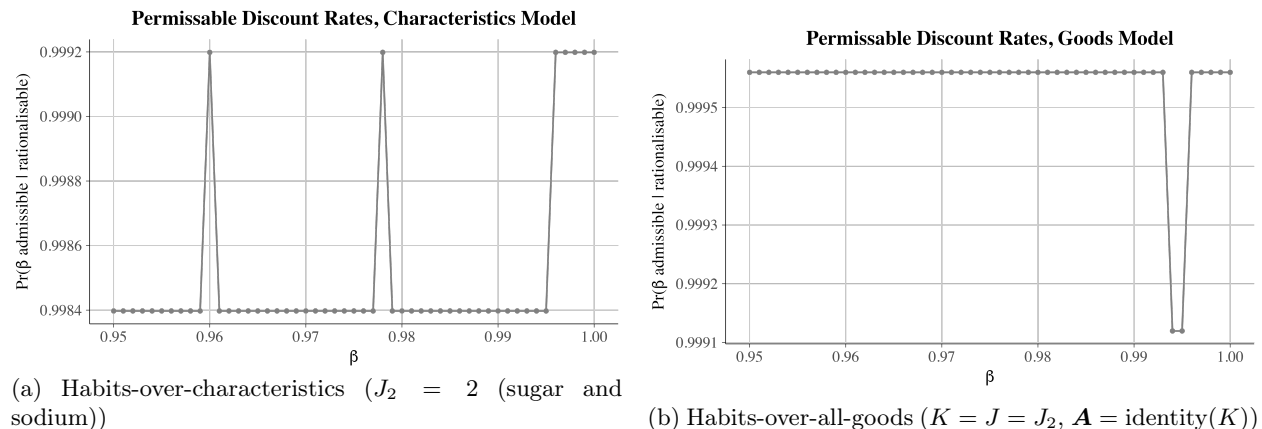
The dynamic RP inequalities are conditional on the discount factor β . To assess informativeness, I evaluate each household over a grid $\beta \in \{0.950, 0.951, \dots, 1.000\}$ and record the set of values for which the inequalities in Theorem 2.2 are feasible. A household is classified as rationalisable if at least one β is admissible.

Figure 9a reports, for each β , the fraction of rationalisable households consistent with that value under the habits-over-characteristics specification. Figure 9b reports the corresponding results under habits-over-goods.

Acceptance probabilities are uniformly high across the grid, with only isolated spikes and dips reflecting discrete feasibility changes for marginal households. There is no systematic monotonic relationship between β and pass rates. In particular, the goods-based model admits nearly the full grid for almost all rationalisable households.

Overall, the identification sets for β are wide. In scanner environments with limited effective price variation across adjacent periods, the dynamic inequalities provide weak discipline on intertemporal discounting once the structural hedonic restrictions are satisfied.

Figure 9: Admissible discount factors among rationalisable households



Notes: For each discount factor β in the grid $[0.95, 1]$, the figure reports the share of households that remain feasible at β , conditional on being rationalisable for at least one value of β . Panel (a) corresponds to the habits-over-characteristics specification with $J_2 = 2$ (sugar and sodium). Panel (b) corresponds to the goods specification with $\mathbf{A} = \text{identity}(K)$.

E Structural and behavioural restrictiveness: additional results

To assess the empirical restrictiveness of each model, I follow the logic of Fudenberg et al. (2023) by evaluating how unusually well the model fits the observed data relative to nearby feasible alternatives. I specify discrepancy measures for the structural and behavioural restrictions implied by the model and compare their realised values in the observed data to those generated by locally perturbed versions of the scanner environment. This comparison yields quantile-based measures of restrictiveness, which record the position of the observed discrepancy within the distribution induced by the eligible set of perturbations.

E.1 Simulation design

The simulation procedure preserves the observed zero pattern in each period, meaning that only goods actually purchased in period t receive positive simulated quantities or prices. This is not a restriction on the underlying choice set but a requirement of the LP characterisation: the equality system ($B2^+$) is defined over the active goods in each period, and altering that active set would change the dimension and geometry of the equality manifold. Preserving zeros therefore fixes the empirical mapping between goods and characteristics, the dimensionality of \tilde{B}_t , and the inputs entering the behavioural inequalities in characteristics space. For each household and model, I generate $M = 10,000$ such locally perturbed datasets.

Within each period, simulated prices for active goods are drawn from a uniform distribution around the household’s observed price range, while simulated quantities are drawn as positive Dirichlet shares and rescaled to match observed total expenditure across periods. This yields a locally perturbed but structurally comparable environment in which both structural and behavioural margins vary across simulations.

E.2 Observed versus simulated moments

Table 11 reports the corresponding discrepancy means in the observed and simulated data. Two features stand out. First, all five characteristics-based models produce identical structural moments, as expected given their common equality manifold, and the observed distances are smaller than those generated under local perturbations (mean 0.167 observed vs. 0.289 simulated). Second, behavioural variation arises only through the CCEI, but these values are near one in both the observed and simulated data. The goods-based models, by contrast, yield distance equal to zero by construction and CCEI essentially equal to one throughout, reflecting their lack of structural and behavioural content.

Table 11: Observed and simulated restrictiveness statistics

Model	Dist (real)	Dist (sim)	CCEI (real)	CCEI (sim)
Habits-over-all-characteristics	0.167	0.289	1.000	0.951
Characteristics (no habits)	0.167	0.289	0.998	0.955
Habits-over-sodium	0.167	0.289	1.000	0.950
Habits-over-sodium-sugar	0.167	0.289	1.000	0.951
Habits-over-sugar	0.167	0.289	1.000	0.956
Habits-over-all-goods	0.000	0.000	1.000	1.000
Goods (no habits)	0.000	0.000	0.998	0.999

Notes: “Dist” denotes the household-level mean structural distance-to-manifold; “CCEI” denotes the household-level behavioural efficiency index. “Real” columns report averages across households in the observed data. “Sim” columns report averages across households and simulation draws under the local perturbation design.

E.3 Quantile-based measures of restrictiveness

For each discrepancy measure, I evaluate the position of the observed data within the distribution induced by locally perturbed environments. For the structural margin, the relevant quantity is the empirical quantile

$$q_{n,m}^{\text{dist}} = \frac{1}{M} \sum_{j=1}^M \mathbf{1}\{d_{n,m,j}^{\text{sim}} \leq d_{n,m}^{\text{obs}}\},$$

which records the fraction of simulated datasets whose prices lie at least as close to the hedonic equality manifold as the observed data. Smaller values indicate that the observed prices are unusually well aligned with the structural restrictions relative to nearby feasible price systems.

For the behavioural margin, I analogously compute the quantile

$$q_{n,m}^{\text{CCEI}} = \frac{1}{M} \sum_{j=1}^M \mathbf{1}\{\text{CCEI}_{n,m,j}^{\text{sim}} \geq \text{CCEI}_{n,m}^{\text{obs}}\},$$

which measures how frequently locally perturbed datasets exhibit behavioural efficiency at least as high as that observed in the data. Smaller values indicate that the observed behaviour is unusually close to dynamic rationality relative to nearby perturbations.

In this application, however, both discrepancy measures exhibit substantial mass at their respective bounds: many households lie exactly on the equality manifold ($d = 0$), and many observed and simulated datasets satisfying the structural equalities attain $\text{CCEI} = 1$. The quantile summaries should therefore be interpreted as measures of relative restrictiveness within a locally perturbed environment, not as simple monotone transforms of the corresponding model-wide mean moments. In particular, the large mean right-tail quantiles for CCEI reflect the prevalence of ties at the behavioural ceiling rather than a contradiction with the lower simulated mean CCEI reported in Table 11.

Table 12: Empirical quantile-based restrictiveness measures

Model	Mean quantile (<i>dist</i>)	$\Pr(q_{n,m}^{\text{dist}} < 0.05)$	Mean quantile (<i>CCEI</i>)	$\Pr(q_{n,m}^{\text{CCEI}} < 0.05)$
Habits-over-all-characteristics	0.579	0.117	0.998	0.00206
Characteristics (no habits)	0.579	0.117	0.990	0.00206
Habits-over-sodium	0.579	0.117	0.998	0.00206
Habits-over-sodium-sugar	0.579	0.117	0.998	0.00206
Habits-over-sugar	0.579	0.117	0.998	0.00206
Habits-over-all-goods	1.000	0.000	1.000	0.000
Goods (no habits)	1.000	0.000	0.994	0.000

Notes: For each household and model, $q_{n,m}^{\text{dist}}$ is the empirical left-tail quantile of the observed structural distance relative to its simulated distribution; $q_{n,m}^{\text{CCEI}}$ is the empirical right-tail quantile of the observed CCEI relative to its simulated distribution. Smaller quantiles indicate greater model restrictiveness. In this application both measures place non-trivial mass at their bounds, so the reported quantiles should be read as relative positions within the local perturbation distribution rather than as direct transformations of the mean discrepancy moments. Entries report means across households and the fraction with quantile below 0.05.

E.4 Interpretation and implications for raw pass rates

The structural quantiles clarify that the hedonic equalities are materially tighter than the goods-based benchmark, but that the observed data are not typically isolated in the extreme lower tail of the local perturbation distribution. This reflects both scanner-price components orthogonal to consumption-relevant characteristics and the fact that a sizeable share of observed and simulated environments lie exactly on the equality manifold.

The behavioural quantiles are less informative in this application because, conditional on satisfying the structural equalities, observed and locally perturbed datasets frequently attain $CCEI = 1$. The large mean right-tail quantiles therefore reflect a ceiling effect rather than evidence that the behavioural restrictions fit the observed data poorly.

The goods-based models achieve raw pass rates above 90% not because they provide a better description of observed behaviour, but because they impose no structural discipline: their distance statistic is identically zero by construction. The quantile-based measures therefore reinforce the view that the main cross-model differences in raw pass rates are structural. Within the class of characteristics-based models, alternative allocations of habit-forming attributes generate nearly identical quantile summaries, so allowing for habits matters mainly at the margin relative to the static characteristics benchmark.