Consumer Choice over Shopping Baskets: A Linear Demand Approach

Afonso Rodrigues*

November 18, 2025

Abstract

Popular demand estimation approaches impose unrealistic choice constraints and misstate market boundaries, biasing price elasticity estimates and affecting our understanding of market power and pass-throughs. I introduce a novel, scalable method that conditions the outcome of consumers' choice on the structure of their *consideration set*: a function of all the combinations of goods – *shopping baskets* – considered each purchase instance by consumers. I show that if consideration sets bind consumers' consumption bundles, joint purchases induce substitutes and complements, possibly making market concentration welfare-increasing. To allow for demand estimation across product categories while addressing dimensionality concerns, I develop a Slutsky matrix proxy from joint-purchase frequencies. I test the model's predictions and jointly determine price elasticities for 20 000 goods across 500 product categories for a Portuguese grocery store sample from 2020-23. The results match observed price volatility, profit margin surveys, as well as reports on shifting consumer tastes during the sample period. Mark-ups are found to have remained volatile around a stable mean, with peaks and troughs corresponding to COVID-related events.

JEL: C14, C36, C38, C55, D12, D83, L81.

Keywords: Discrete choice, continuous choice, consumption bundles, joint purchases, shopping baskets, market power, consideration set, complements, substitutes, dimensionality reduction, demand estimation, point-of-sale data, quadratic utility, linear demand.

---

*Rodrigues: Department of Economics, University of Oxford (afonso.rodrigues@economics.ox.ac.uk). I have no relevant or material financial interests that relate to the research described in this paper. I gratefully acknowledge financial support by the Foundation for Science and Technology of the Ministry of Education, Science and Innovation (Portugal), as well as data sharing by SONAE. The data agreement involves a request for review of the findings prior to their release. I am indebted to my dissertation advisor Alexei Parakhonyak, as well as Ian Crawford, Alex Teytelboym, Howard Smith, Abi Adams, and Giulio Gottardo for their helpful insights. I benefited from thoughtful conference discussions by Leonardo Madio and Anne-Christine Barthel, as well as valuable comments from participants of the NIE PhD Symposium 2024, the 13th Oligo Workshop, the Research Jamboree at the University of Oxford, the Lisbon Meetings in Game Theory and Applications #14, and EARIE 2025.

1 Introduction

When accurate, demand estimation provides a picture of how buyers react to prices and product features, allowing us to say something concrete about market power, assess mergers, and track spillovers. It also lets us gauge the effects of taxes, subsidies, or new rules before they are put in place, allowing for counterfactual analysis of options and trade-offs.

In the standard model of choice behaviour in Economics, choice is characterised as the result of perception, driving preferences, converging into a process which produces a choice (McFadden, 2001). Perception and preferences are defined over goods and their attributes; goods are also the available choices. Textbook rational consumers then consider every feasible consumption bundle of these goods, selecting the combination which yields the greatest utility affordably (see e.g. Mas-Colell et al., 1995).

This base setup raises two problems for demand estimation. First, the assumption that consumers consider every feasible consumption bundle is unrealistic. For example, guacamole is a popular spread, yet consumers hardly ever buy avocados and red onions without limes; we may assume that they only consider either the combination or the individual goods, but not the pairs. This also challenges the idea of preferences and choices over goods. The choice to buy the combination of avocados, red onions, and limes is lower-dimensional than their separate choice - i.e. demand for each good is dependent on that of the others. This dependence may have implications for the outcomes of consumer choice.

Second, even when adequate, standard assumptions remain computationally infeasible. Gandhi and Nevo (2021) discuss at length the core challenges of demand modelling and estimation for differentiated goods. In practice, empirical applications target discrete, individual-good choice for a few goods within a product category. This approach can be extended to account for inattention, bounded rationality, search costs, and (unobserved) constraints (see e.g., Abaluck and Adams, 2021). Yet discrete-choice models are restrictive in their own way. The allowable set of consumption bundles is simultaneously too small and too large: many cross-category interaction effects are left out - e.g. butter and bread, milk and cereal; and many goods are included despite being rarely considered by themselves - such as add-on sales, impulse purchases, non-food items at grocery stores. As a whole, this underestimates price elasticities, producing excessive estimates of market power and understating pass-throughs.

In this paper, I introduce an approach which addresses the computational challenges of allowing for broad flexibility in consumer choice while imposing realistic constraints. I employ a simple model of linear demand by a representative consumer, micro-founded from quasi-linear quadratic utility. Rather than deciding which goods to purchase, the representative consumer conditions the outcome of their utility maximisation problem on the structure of their *consideration set*, and selects amongst the alternatives in it.

This structure effectively controls the space of possible consumption bundles that the utility function can produce. Think of all the combinations of goods that a consumer considers each purchase instance. Consumption bundles are determined in practice by the consumer's choice process across all these small moments. To shape how choice is restricted by consideration set structure, I therefore define the latter as a function of all the combinations of goods considered each purchase instance by consumers. Shopping baskets may also include goods purchased standalone. The consideration set is assumed exogenous and stable. Instead of micro-founding the consideration set, I use this assumption to generalise, and study what consideration set characteristics constrain consumer choice, via conditional outcomes.

This yields novel theoretical predictions. I show that the consideration set binds if and only if some goods are never considered for purchase by themselves. This is because goods consumption must be a non-negative, linear function of the set of shopping baskets considered by consumers (NN-LF conditions hereafter); every good not considered individually restricts the functions the consideration set can produce irredeemably. If the consideration set binds, joint purchases induce substitutes and complements onto the system. When the structure of the consideration set binds, own-price effects are weaker than they would otherwise be, as goods' demand becomes relatively more dependent on those goods with which they are often co-purchased. Often co-considered goods behave as complements, while goods often paired with the same products face substitution effects conditional on these. The constrained choice behaviour imposes lower average demand, aggregate profits, and consumer surplus. A major implication of the outsized role of consideration is that product ownership concentration may become welfare-enhancing.

Incorporating these predictions onto demand estimation enables modelling under more realistic conditions. To estimate linear demand, I use point-of-sale data from across a sample of grocery stores owned by a supermarket chain in Portugal, for the period from February 2020 to February 2023. The data covers more than 20 000 goods and 500 product categories. The representative consumer's consideration set is derived from the set of unique transactions observed during this period. To address the Gandhi and Nevo (2021) critique, I expand on the Pinkse and Slade (2004) semi-parametric approach, proxying for the Slutsky matrix of the demand system via a statistical model of joint-purchase frequencies (Tian, Lautz, Wallis, and Lambiotte, 2021).

There are many reasons why consumers may jointly purchase goods, which the literature splits into two broad types: preferences and correlations (see Gentzkow, 2007, for an early formulation of this). Correlated purchases can be the outcome of things like minimising trip-related costs or time and stock considerations. Preferences toward co-purchases meanwhile may come from a love for variety or complementarities.

To isolate the latter, the methodology in Tian et al. (2021) works as follows. The data is organised into raw counts of how often every pair of items is jointly purchased, and we note both the number of products per transaction and the number of transactions per product. These are then used to adjust for product popularity and shopping basket size, to minimise the influence of correlations. A statistical "null" model – called a Bipartite Configuration Model - asks: "Given how popular each item is and how large shopping baskets typically are, how often would two items show up together just by chance?" Complements are co-purchased significantly more often than expected, substitutes share complements and are co-purchased significantly less often than expected. The derived substitutes matrix, a product-by-product mapping, is fed into the demand model via polynomial series expansion to predict price effects across products. Different order polynomials allow flexibility to observe both substitutes and complements; they mimic spillover/lagged effects in spatial and network models, with every polynomial parameterised to assess how products are influenced by their competitors, competitors-of-competitors, etc.

The empirical model is Two-Stage Least Squares, with a set of linear demand functions, one per good with supermarket and time FEs. Product utility is a function of observed product characteristics, including product category and private-label brand dummies. Prices are interacted with these to allow for heterogeneity. I use a single exogenous instrumental variable – the average price across each goods' competitors at other grocery stores (Hausman, 1996). The polynomial lags of the Slutsky matrix proxy are interacted with the instrument in the same way as they are with the price vector, allowing a single instrument to cover every price parameter. All variables are standardised for comparability, and I orthogonalise the set of covariates and price interactions to eliminate multicollinearity and reduce the number of parameters to estimate.

The NN-LF conditions, when binding, affect linear demand estimation in two ways. First, the specific Slutsky matrix that the proxy is meant to approximate changes, as established by the abovementioned theoretical findings. Second, the non-negativity condition on the alternatives in the consideration set introduces a form of price endogeneity in the system. Ignorance of such a constraint causes prices to be correlated with the error term because, as prices rise, the necessary adjustment to force negative demands back to zero also rises. The estimation of demand with binding non-negativity constraints is not new (Wales and Woodland, 1983); but here the non-negativity is not imposed at the product, but at the consideration set alternative level. To address this, I use a novel control-function correction with a non-negative least-squares step.

In doing so, I successfully recover price elasticity parameters, which are used to measure mark-ups across the product assortment. I construct a revenue-weighted price index and document pronounced volatility aligned with Portugal's four COVID-19 waves, with price peaks lagging each wave; prices broadly co-move with the CPI but decouple late in the sample, as inflation in the energy and logistics sectors takes off in 2022. Between 2020 and 2023, mark-ups remain stable across all major product categories. I document revenue-weighted mean mark-ups averaging $\approx 11\%$, with a trough in the first quarter of 2021 ($\approx 8.5\%$) and a peak in the third quarter of 2022 ($\approx 15\%$) - levels substantially below those previously documented (e.g. Dopper, MacKay, Miller, and Stiebale, 2024). Benchmarking against U.S. grocery survey margins on a four-quarter moving average basis reveals near-identical timing and volatility with observed revenue-weight prices and mark-ups, and suggest strong pass-through, consistent with my estimate of a low-mark-up environment; level differences are also consistent with store margins inclusive of upstream supplier mark-ups.

Robust measures of mark-up distribution over the sample period show that changes are generally insignificant except for an $\approx 10\%$ rise among top-percentile products beginning in the first quarter of 2022, and persisting two quarters, coincident with the CPI–price-index decoupling. Importantly, the 2022 increase is not driven by supermarket heterogeneity. To study the drivers of this late mark-up rise, I study transitions between mark-up percentiles, finding high persistence with mild mean reversion, but weakened persistence in mid-to-late 2022. This is consistent with shifting tastes and emergence of new high-mark-up goods. Analysis of fixed-basket Laspeyres- and Paasche-style mark-up indices further indicate reallocation, with the Paasche exceeding the Laspeyres by the third quarter of 2022.

2 Literature review

In empirical Industrial Organisation, early attempts to estimate choice behaviour (e.g. Theil, 1965; Christensen et al., 1975; Deaton and Muellbauer, 1980) aimed to allow demand estimation across large numbers of goods flexibly, but faced some difficulties (Gandhi and Nevo, 2021). In these models, as the number of alternatives grows larger, the number of own- and cross-price effect parameters grows (proportional to) quadratically in these, producing a curse of dimensionality. Price endogeneity also becomes nearly impossible to address, due to collinearity and the need to find enough exogenous (and non-collinear) instruments to separately identify price effects. Other issues, such as addressing consumer heterogeneity and the entry of new goods, are also relevant difficulties.

Constant Elasticity of Substitution (CES) demand (Spence, 1976; Dixit and Stiglitz, 1977) and the logit model (introduced by McFadden, 1973) offered responses to these limitations, with significantly different utility functional forms and assumptions. The latter has become especially popular in empirical Industrial Organisation; Berry, Levinsohn, and Pakes (1995) remains the key reference for demand estimation via logit/discrete-choice models. Other significant contributions in this tradition include Nevo (2001) for supermarket scanner data; and Hausman, Leonard, and Zona (1994) on instrumental variables for price endogeneity. Discrete-choice models enable modelling individual choice under consumer heterogeneity and allow for correlations in unobserved factors. However, their fit for applications over large choice sets is limited due to a *curse of dimensionality*, where the number of parameters to be estimated in the correlation matrix is always more than proportional to the number of goods, and therefore mixed logit models are far too large to be manageable for a large number of products. Compounded by computational difficulties (see, e.g., Lanier, Large, and Quah, 2023) and concerns about large market asymptotics (Armstrong, 2016), this remains a work in progress.

Researchers' ability to study competition and market structure in settings with large choice sets has thus been limited. Dopper, MacKay, Miller, and Stiebale (2024) use scanner data across over 100 consumer product categories to estimate demand and recover mark-up trends. The authors apply a random coefficient logit demand system and a covariance-restrictions approach to each individual category and estimate demand for each in isolation. Fosgerau, Monardo, and de Palma (2024) achieve further scalability and flexibility in a similar setting, employing a generalised inverse logit model, optimised for computational efficiency by leveraging market segmentation. Their approach effectively allows for a much greater number of alternatives than previous versions of logit model by adjusting the model such that it can be estimated via linear instrumental variables regression with market-level data. This linear approach echoes that of the present paper. However, it does not explicitly overcome the concerns regarding large market asymptotics (Armstrong, 2016) which discrete choice raises, nor does it directly account for consumers considering combinations of goods over individual product purchases.

The present paper instead builds on the earlier attempts. In particular, we pursue an alternative framework based on consideration sets for thinking about multi-good cross-category purchases. This differs from the multi-level demand systems of Gorman (1959), Deaton and Muellbauer (1980), and Hausman, Leonard, and Zona (1994), where the weak separability assumption in the utility function blocks most cross-group interactions, except via the top budget. I adopt a tractable linear-demand framework with a representative consumer, derived from a quasi-linear quadratic utility (see Amir, Erickson, and Jin, 2017, for a deep dive; see Choné and Linnemer, 2020, for a historical retrospective). Linear demand remains the workhorse model in Economics textbooks due to straightforward intuition and application, closed-form solutions, and parsimony. The representative agent conditions the solution to their utility-maximisation problem on the architecture of a fixed consideration set. Rather than choosing goods directly, the agent selects among the alternatives contained in that set; these alternatives, or shopping baskets, can be single items or combinations. This arrangement effectively delimits the feasible set of consumption bundles that the utility specification can generate. Rather than micro-founding the consideration set itself, I treat the consideration set as exogenous and time-invariant to generalise and study how consideration sets constrain choice and condition outcomes. Rather than focusing on a single category or pair of categories, the present paper also attempts to estimate demand across the entire assortment of goods present in the dataset (for a previous attempt, see Thomassen, Smith, Seiler, and Schiraldi, 2017).

Consideration sets, first established in Wright and Barbour (1977), are a relaxation of the assumption that individuals consider all feasible choices. They can be defined as the restricted set of shopping baskets to which a consumer pays attention in their choice process. Many reasons have been given as to why the choice set (the set of available choices) may sometimes differ from a consumer's consideration set: behavioural heuristics (Hauser, 2014), cognitive limitations (Simon, 1957), bounded rationality and seller persuasion (Eliaz and Spiegler, 2011), imperfect information (Goeree, 2008), search (Caplin, Dean, and Martin, 2011), or evaluation costs broadly speaking (Hauser and Wernerfelt, 1990). Consideration sets have most recently been applied to discrete-choice models (Matějka and MacKay, 2015) and in models of revealed preference (Abaluck and Adams, 2021). Another recent contribution, Crawford, Griffith, and Iaria (2021) surveys how assumptions over the structure of consideration sets can impact demand estimation.

One concern the present paper speaks to is the adjustment of consideration sets to settings where it is probable that consumers consider goods both individually and in combination (e.g.

Manchanda, Ansari, and Gupta, 1999; Fox and Lazzati, 2017; Allen and Rehbeck, 2022). Iaria and Wang (2019) develop a structural demand model to analyse how consumers choose product bundles, allowing for both complementarity and substitutability between goods. Their mixed logit framework incorporates a novel demand inversion technique that reduces some of the aforementioned computational complexity. Applying their model to the same data and number of brands as in Nevo (2001), they find strong evidence of demand synergies across products, driving joint purchases within the same transaction. They demonstrate that ignoring these interactions leads to biased estimates and misleading policy conclusions. Synergies were first discussed in Gentzkow (2007). Ershov, Laliberté, Marcoux, and Orr (2024) make a similar methodological contribution also built on a random-coefficients logit framework. They develop a Generalised Method of Moments (GMM) estimator that uses an aggregation structure based on preference restrictions to combine aggregate sales data with specific choice probabilities per unique combination of soda, chips, or soda and chips. This results in a contraction mapping for the three aggregate "inside" goods. They address price and characteristic endogeneity using standard instrumental variables and a set of novel shift-share instruments for product variety. In an empirical review of the U.S. chips and soda market using scanner data, the authors are able to study an average of over 1300 soda and 800 chip varieties across all market-year pairs. They find large cross-category complementarities, suggesting that mergers between multi-product multi-category firms lead to lower price increases and even price decreases compared to a substitution-only model. However, this approach allows for only a small number of product categories (two in their case).

Two recent working papers come even closer to this paper's approach. Lewbel and Nesheim (2019) build a demand model in which consumers choose a few goods and quantities, allowing substitutes and complements with corners and random coefficients, and show semiparametric identification. In Lewbel and Nesheim (2019), all possible combinations of goods are considered by consumers, which is likely unrealistic. Their estimation procedure, reliant on demand sparsity, only mitigates the resulting curse of dimensionality, even while their model is, like this paper's, also rooted in quadratic utility. Nocke and Schutz (2025) take a potential-games approach, introducing transformed potentials and characterising IIA/generalised-linear and nest/basket demand classes that permit complementarities and guarantee equilibrium existence under multi-product firm-pricing games. Nocke and Schutz (2025) thus provide a complementary theoretical characterisation to the present paper, though one with an also unrestricted consideration set. Nocke and Schutz (2025) find that whether products are substitutes or complements depends on the interaction effects between goods and outside option but also through joint purchases - the extent to which the products of two different firms are predominantly purchased together or not determines the strength of complementarity. The present paper complements these by giving an empirically scalable, consideration-set–based way to map joint purchases into cross-price effects and market-level elasticities.

My semi-parametric demand estimation method is closest to that of Pinkse and Slade (2004). Their "spatial differentiation" approach proxies the latent Slutsky matrix of the demand system via series expansion of a distance-based measure of product differentiation. By semi-parameterising the cross-price effects, sufficient variation is introduced in the model for identification - with further heterogeneous own-price effects via interactions with covariates. This structure effectively regularises the system and captures latent complementarity and substitutability patterns that (1) help de-noise and disentangle otherwise collinear price effects, and (2) eliminate dimensionality concerns.

My proposed approach is behaviourally and data richer, as well as more computationally robust. Besides the inclusion of consideration set constraints, to minimise collinearity which arises from the Pinkse and Slade (2004) approach, the series expansion terms interacted with prices and the covariates are orthogonalised via principal components analysis (PCA). This helps reduce the number of explanatory variables (minimising overfitting) and does away with collinearity concerns entirely, minimising bias. The distance-based measure of product differentiation in this paper – a statistical model of co-purchase frequency first proposed in Tian et al. (2021) - is directly derived from observed purchase behaviour, allowing for greater granularity of substitute and complementary relationships between goods as understood by consumers.

Distance-based measures of product differentiation themselves are rooted in the earliest papers in Industrial Organisation theory. Lancaster (1966)'s product characteristics model, the Hotelling (1929) line, or the Salop (1979) circle are all core papers in the field, which view product differentiation as a matter of location in an abstract product space. More recently, Hoberg and Phillips (2016) have popularised cosine similarity, taking advantage of text-based measures of product similarity to identify competitors. They target a subset of firms/goods, define them over a set of descriptives identified via textual analysis of U.S. public firms' 10-K filings, and compute the cosine similarity between these sets. They assume this cosine similarity approximates the similarity between the goods' real attributes; this approach is thereby said to situate these goods in a multi-dimensional product space. In this space, the cosine similarity defines the angle between two goods with said angle assumed to negatively correlate with degree of competitive pressure. Readers are encouraged to see Ushchev and Zenou (2018), which generalise the theory behind this approach. Distance-based measures are also popular sources of instrumental variables (Gandhi and Houde, 2019).

Other papers have recently experimented with semi-parametric estimation of price effects, similarly conceptualising product differentiation via distance-based measures and applying these empirically (e.g. Dubois, Griffith, and Nevo, 2014; Magnolfi, McClure, and Sorensen, 2025; Pellegrino, 2025). For these papers, the main issue is the data required to produce the cross-price effects matrix. Dubois et al. (2014) relies on characteristics data, which is simultaneously hard to obtain and gives the researcher excessive oversight over which characteristics matter and how to measure them; Pellegrino (2025)'s approach only works for publicly listed, (near) single-product firms; and Magnolfi et al. (2025) presents methodological difficulties in survey scaling. In using widely available co-purchase data, these limitations can be surpassed.

3 Illustrative example

For any set of K goods, there are up to 2^K unique feasible combinations. In many settings, especially those with high K, it may therefore be unrealistic to assume consumers consider for purchase all feasible combinations.

For illustration purposes, consider the following setting. A representative consumer buys their groceries at a shop that only offers three goods: milk, bacon, and pasta. They only ever go to the shop if they wish to buy: breakfast, in which case they invariably choose to buy milk and bacon; dinner, when they strictly eat bacon and pasta; or when they go do their weekly shopping, which, depending on their mode of transport, may imply purchasing greater or lower quantities of all three goods. In other words, we may express our representative consumer's consideration set as follows:

$$\underset{\times choices}{A_{\{milk,bacon,pasta\}}} = \begin{pmatrix} 2 & 0 & 1 & 2 \\ 2 & 2 & 2 & 4 \\ 0 & 2 & 1 & 2 \end{pmatrix} \tag{1}$$

From this mapping between goods and choices, we can then express goods consumption bundle, $\boldsymbol{q}_{K\times1}$, as a function of how many times each of the shopping baskets present in the consideration set was selected across all purchase instances, $\boldsymbol{z}_{J\times1}$:

$$\boldsymbol{q} = A \cdot \boldsymbol{z} \tag{2}$$

Both $\boldsymbol{q}$ and $\boldsymbol{z}$, as well as A, are non-negative. Note that this approach nests the case where the consumer chooses over individual goods: $\boldsymbol{q} = I \cdot \boldsymbol{q}$, for I the identity matrix. This is equivalent to a setting where there are no consideration set restrictions on the consumption bundle: quasi-linear quadratic utility then allows any $\boldsymbol{q} \in \mathbb{R}^K_+$.

In quasi-linear quadratic utility models (QQUMs), a representative consumer[1]'s utility maximisation problem is given by:

$$\max_{\boldsymbol{q}} \; U(\boldsymbol{q}) = \boldsymbol{q}'\boldsymbol{\delta} - \frac{1}{2}\boldsymbol{q}'M\boldsymbol{q} + q_0 \quad \text{s.t.} \quad \boldsymbol{q}'\boldsymbol{p} + q_0 \leq Income \quad \& \quad \boldsymbol{q} \geq 0 \tag{3}$$

$\mathcal{K}$ is the set of all differentiated goods $k = 1, \ldots, K$ said consumer optimises its consumption over, with $\boldsymbol{q}$ the $K \times 1$ vector of consumption per good. The parameter $\boldsymbol{\delta}$ is a strictly positive $K \times 1$ vector of marginal product utilities. The representative consumer draws utility from each product both directly and via interaction effects across those it purchases. These interaction effects between products are defined by a $K \times K$ sparse matrix M. The off-diagonal entries of the matrix may have any sign. For example, if $M_{12} > 0$, then consuming a greater amount of good 1 lowers the additional amount desired of both that good, due to satiation, and of 2. Lastly, q_0 is the numeraire.

[1]For a discussion of the representative consumer assumption, see the Appendix.

The properties of the QQUM are well-understood.[2] M must be positive definite, as to satisfy (i) strict concavity, (ii) downward-sloping demand, and (iii) dominant own-price effects. Despite this, the M matrix can accommodate rich patterns of substitution and complementarity. The model assumes no income effects.[3] The representative consumer objective function can be re-written as follows:

$$\mathcal{L}(\boldsymbol{q}) = \boldsymbol{q}'(\boldsymbol{\delta} + \phi\boldsymbol{p}) - \frac{1}{2}\boldsymbol{q}'M\boldsymbol{q} + q_0 \quad \text{s.t.} \quad \boldsymbol{q} \geq 0 \tag{4}$$

The (negative) term ϕ maps price sensitivity/marginal utility of income. Initial marginal utilities are assumed to be positive, i.e., $\boldsymbol{\delta} + \phi\boldsymbol{p} > 0$.

Quasi-linearity allows the derived demands to depend solely on the vector of prices (without assumptions about market size or income effects) and the interactions among goods, as captured by an interaction matrix. In effect, the QQUM yields linear demand functions of the following form:

$$\boldsymbol{q}_g = M^{-1}(\delta + \phi\boldsymbol{p}) \quad \text{s.t.} \quad \boldsymbol{q}_g \geq 0 \tag{5}$$

with price elasticity proportional to the inverse of the interactions matrix:

$$\frac{\partial q_i}{\partial p_l} = \phi\{M^{-1}\}_{il}, \ \forall i, l \in \mathcal{K} \tag{6}$$

Consider now what happens if we instead assume that consumer choices depend on the consideration set we have defined above. Through the identity above, we can re-write the representative consumer's utility maximisation problem as:

$$\max_{\boldsymbol{z}} \ U(\boldsymbol{q}) = \boldsymbol{q}'\boldsymbol{\delta} - \frac{1}{2}\boldsymbol{q}'M\boldsymbol{q} + q_0 \quad \text{s.t.} \quad \boldsymbol{q}'\boldsymbol{p} + q_0 \leq \mathit{Income} \quad \& \quad \boldsymbol{q} = A \cdot \boldsymbol{z} \quad \& \quad \boldsymbol{z} \geq 0 \tag{7}$$

Three key changes have been made. Utility is now maximised via choice of $\boldsymbol{z}$; $\boldsymbol{q}$ is a function of $\boldsymbol{z}$, and there is now a non-negativity constraint on $\boldsymbol{z}$ rather than on $\boldsymbol{q}$ (though the latter follows implicitly from the new conditions). As I detail later, optimising in $\boldsymbol{z}$ we obtain the following demand function:

$$\boldsymbol{q}_b = A(A'MA)^{+}A'(\boldsymbol{\delta} + \phi\boldsymbol{p}) \quad \text{s.t.} \quad \boldsymbol{z} \geq 0 \tag{8}$$

for $(\cdot)^{+}$ the Moore-Penrose pseudo-inverse. Assuming e.g. $M = 0.1 \times \mathbf{11'} + (1 - 0.1)I$, $\boldsymbol{\delta} = [2 \ 2 \ 2]'$, $\phi = -0.1$, and $\boldsymbol{p} = [1 \ 1 \ 1]'$:

$$\boldsymbol{q}_g = [1.58 \ 1.58 \ 1.58]' \neq [1.09 \ 2.17 \ 1.09]' = \boldsymbol{q}_b \tag{9}$$

[2]See Amir, Erickson, and Jin (2017) and Choné and Linnemer (2020) for an exhaustive treatment.

[3]This may be restrictive in some scenarios. The reader is directed to Tirole (Introduction, p.11, 1988) and Vives (1987) for the relevant conditions.

In other words, which consumer choice model we choose may yield substantially different outcomes, even in a relatively straightforward setup.

There can be many reasons why consumers may display behaviour compatible with the existence of a consideration set. Some shopping baskets may tend to include products that "go together" because they serve a common need: a desire to buy breakfast, dinner, or go on the weekly shopping run, as above; seasonality or promotions, where a shopping basket might reflect seasonal trends (holiday specials, summer outdoor items) or be influenced by a particular promotion; and category or style, where purchases might cluster by category (electronics, clothing, groceries) or even by style (premium vs. budget items). Turns out, just the fact that a consumer considers more goods across choices than the total number of choices considered can be enough to obtain the outcome above. More on this below.

Having demonstrated the pertinence of evaluating the causes of the distinct outcomes observed, let us proceed to formalise a consideration-set-based theory of consumer choice.

4 Setup

Assume that a representative consumer restricts its purchase decisions to a fixed, exogenous consideration set. A consideration set is the subset of a full set of feasible alternatives, that a consumer pays attention to and evaluates when making a decision. Let $\mathcal{J}$ be the discrete set of these alternatives, or shopping baskets, $j = 1, \ldots, J$ considered by a representative consumer, and $A_{K\times J}$ the matrix representation of the consideration set: a sparse matrix of product presence per unique basket considered, each element a_{kj} constituting the number of units of a given good k in a given basket j considered. Each column can be thought of as the equivalent to one possible receipt during a purchase instance. Vector $\boldsymbol{z}_{J\times 1}$ is the consumer's level of consumption per basket (divisible for simplicity). As written above, we may express the goods consumption bundle $\boldsymbol{q}$ as a *non-negative linear function* of $\boldsymbol{z}$:

$$\boldsymbol{q} = A\boldsymbol{z} \qquad \& \qquad \boldsymbol{z} \geq 0 \tag{NN-LF}$$

This expression (henceforth the NN-LF condition) is evocative of Lancasterian product characteristics models, with the particularity of $\boldsymbol{q}$ being on the LHS and the choice variable on the RHS. These differences evolve into an entirely different set of mathematical calculations as a result, and have different implications.

The NN-LF condition summarises the constraints imposed by a representative consumer's consideration set on choice. When may NN-LF bite? When a consumer makes purchase decisions, the set of all feasible goods consumption bundles will span the *conical hull* of matrix A:

$$\boldsymbol{q} \in cone(A) = \Big\{ \sum_{j=1}^{J} \boldsymbol{a}_j x_j : \boldsymbol{a}_j = A\boldsymbol{e}_j, x_j \in \mathbb{R}_+^K, j \in \mathcal{J} \Big\} \tag{10}$$

Vector $\boldsymbol{e}_j$ is the j-th standard basis vector. In standard consumer choice, $\boldsymbol{q} \in \mathbb{R}_+^K$, so the constraints will play a role depending on the answer to the following question: when is

$cone(A) \subsetneq \mathbb{R}_+^K$?

Proposition 1*: Rank*$(A) < K$ *is a sufficient condition for the consumption bundle* $\boldsymbol{q}$ *to be confined to* $cone(A) \subsetneq \mathbb{R}_+^K$*. The columns of* A *not generating all standard basis vectors* $\boldsymbol{e}_1, \ldots, \boldsymbol{e}_K$ *via non-negative combinations is a necessary and sufficient condition for the consumption bundle* $\boldsymbol{q}$ *to be confined to* $cone(A) \subsetneq \mathbb{R}_+^K$*.*

The first part of **Proposition 1** follows directly from the linear dependence of matrix A. It is a helpful detail which enables us to isolate the $\boldsymbol{q} = A\boldsymbol{z}$, linear function (LF) condition from the broader NN-LF. The second part of the proposition can be proved as follows:

Proof: Suppose $cone(A) \subsetneq \mathbb{R}_+^K$. Then there exists at least one vector $\boldsymbol{x} \in \mathbb{R}_+^K$ such that $\boldsymbol{x} \notin cone(A)$. In particular, consider the standard basis vectors $\boldsymbol{e}_1, \ldots, \boldsymbol{e}_K$. If all of these vectors were in $cone(A)$, then by taking nonnegative linear combinations, we could generate any vector in $\mathbb{R}_+^K$, implying $cone(A) = \mathbb{R}_+^K$, a contradiction; $cone(A) \subsetneq \mathbb{R}_+^K$ implies that at least one $\boldsymbol{e}_k \notin cone(A)$. Conversely, suppose there exists some k such that $\boldsymbol{e}_k \notin cone(A)$. Then, the cone cannot generate all vectors in $\mathbb{R}_+^K$, because it fails to generate the direction corresponding to good k alone. Thus, $cone(A)$ must be a strict subset: $cone(A) \subsetneq \mathbb{R}_+^K$. $\square$

In other words, the NN-LF condition holds if and only if not all goods are considered as standalone alternatives.

Both the rows and columns of matrix A are unique and non-zero, but may not always be linearly independent. In fact, if matrix A has more rows than columns, that is to say, if the representative consumer considers more products across alternatives than the total number of alternatives it considers, then linear dependency is assured (Hurewicz and Wallman, 1941). Even if instead there are more columns than rows, one can think of conditions under which linear dependence may still be produced. If the purchase patterns are driven by a limited number r of latent factors (the "common needs" that cause certain products to be bought together) and if r is smaller than the number of products, then the rows of A will lie in an r–dimensional subspace, and we will once again have linear dependency.

Thus, in general, $\boldsymbol{q}$ may not take arbitrary values in $\mathbb{R}_+^K$; utility maximisation over baskets limits the range of product quantities demanded which we may observe. As shown earlier, this may have significant implications for equilibrium outcomes; our illustrative example is one such case where $rank(A) < K$. In the example, row 2 of $A_{\{milk,bacon,pasta\}\times choices}$ is equal to the sum of rows 1 and 3.

How may we obtain the demand function? Through the NN-LF condition, we can re-write the representative consumer objective function as:

$$\mathcal{L}(\boldsymbol{z}) = \boldsymbol{z}'A'(\boldsymbol{\delta} + \phi\boldsymbol{p}) - \frac{1}{2}\boldsymbol{z}'A'MA\boldsymbol{z} + z_0 \quad \text{s.t.} \quad \boldsymbol{z} \geq 0 \tag{11}$$

Note that $q_0 = z_0$ w.l.o.g. Maximising the Lagrangian of the representative consumer's basket choice problem in $\boldsymbol{z}$, we obtain the following FOC:

$$(A'MA)\boldsymbol{z} = A'(\boldsymbol{\delta} + \phi\boldsymbol{p}) \tag{12}$$

Expressing (12) in terms of $\boldsymbol{z}$ introduces the first problem derived from our dimensionality constraint. Unlike M,

***Lemma 1**: $A'MA$ is positive semi-definite, and therefore not generally invertible.*[4]

To proceed, I multiply both sides of the FOC by $(A'MA)^+$, for $(\cdot)^+$ the Moore-Penrose pseudoinverse (Moore, 1920; Penrose, 1955). However, on the LHS, $(A'MA)^+(A'MA) \neq I$ does not hold in general. A general form of (12) relative to $\boldsymbol{z}$ must account for the so-called *homogeneous part* of differential equations:

$$\boldsymbol{z} = (A'MA)^+A'(\boldsymbol{\delta} + \phi\boldsymbol{p}) + (I - (A'MA)^+(A'MA))\boldsymbol{y} \quad \& \quad \boldsymbol{z} \geq 0 \tag{13}$$

for $\boldsymbol{y}$ a random vector.[5] There are thus a multiplicity of mappings of $\boldsymbol{p}$, as the strict concavity condition we had imposed on M no longer holds for $A'MA$. Nonetheless,

***Proposition 2**: While the mapping from prices to the consideration-set-level choices $\boldsymbol{z}$ is non-unique, the resulting product demand $\boldsymbol{q} = A\boldsymbol{z}$ is unique:*

$$\boldsymbol{q} = A(A'MA)^+A'(\boldsymbol{\delta} + \phi\boldsymbol{p}) \quad \& \quad \boldsymbol{z} \geq 0 \tag{14}$$

Proof: If $\boldsymbol{q} = A\boldsymbol{z}$, $A(I - (A'MA)^+(A'MA))\boldsymbol{y} = 0$, as $I - (A'MA)^+(A'MA)$ is in the null space of $A'MA$, which itself matches the null space of A: $\boldsymbol{x}'A'MA\boldsymbol{x} = (A\boldsymbol{x})'M(A\boldsymbol{x}) = 0$ if $A\boldsymbol{x} = 0$. Since our M matrix is assumed positive definite, the only way for us to have $(A\boldsymbol{x})'M(A\boldsymbol{x}) = 0$ is precisely if $A\boldsymbol{x} = 0$. □

In short: despite the NN-LF condition breaking the strict concavity assumption generally required for the well-behavedness of the utility function, its ensuing use - following the pseudoinverse and resulting non-unique $\boldsymbol{z}$ - ultimately isolates the single unique mapping from $\boldsymbol{p}$ to $\boldsymbol{z}$ needed for obtaining $\boldsymbol{q}$.

The sparsity of the A and M matrices minimises individual product demand's reliance on the price vector of the entire product assortment (and vice-versa). Calculating the cross-price elasticities of demand - abstaining for the time being from the NN condition $\boldsymbol{z} \geq 0$ - the remaining effects are held to be proportional to the expression $A(A'MA)^+A'$:

$$\frac{\partial log(q_a)}{\partial log(p_b)} = \phi \frac{p_b}{q_a} \boldsymbol{a_a}'(A'MA)^+\boldsymbol{a_b} \lesseqgtr 0, \ \forall a, b \in \mathcal{K} \tag{15}$$

Despite losing positive definiteness (and nice properties such as diagonal dominance) under the dimensionality constraint, the demand system is still well-behaved. To see this, note:

[4] Proofs of all Lemmas may be found in the Appendix.

[5] One may verify the propriety of this operation by noting that $(A'MA)(I - (A'MA)^+(A'MA)) = 0$, and $(A'MA)(A'MA)^+A'(\boldsymbol{\delta} + \phi\boldsymbol{p}) = A'(\boldsymbol{\delta} + \phi\boldsymbol{p})$. The latter follows from the properties of the pseudo-inverse, as both $A'MA$ and $A'(\boldsymbol{\delta} + \phi\boldsymbol{p})$ are in the column space of A'.

Lemma 2: *$A(A'MA)^{+}A'$ is symmetric and positive semi-definite.*

Symmetry and positive semi-definiteness means (14) still satisfies the Law of Demand[6] and features equilibrium-dependent heterogeneous elasticities of substitution. Regardless of whether goods are complements or substitutes, the total cross-price effect will be determined both by direct interaction effects between products, but also how the remaining products across the assortment interact with these in the consideration set. This matches intuition: demand for a given good will be driven also by the substitutes and complements of those goods with which it is most often considered and purchased.

5 Equilibrium

Assume first that all K differentiated products in the assortment are sold by K different single-product firms. Firms compete toward the Bertrand (Nash) equilibrium. For simplicity of exposition, marginal costs are set to $c_i = 0, \forall i \in \mathcal{K}$. The representative consumer takes prices as given.

Proposition 3: *When each of the K differentiated products is sold by a separate firm engaging in Bertrand (Nash) competition, with the representative consumer's demand given by $\boldsymbol{q} = A(A'MA)^{+}A'(\boldsymbol{\delta} + \phi\boldsymbol{p}) \quad \& \quad \boldsymbol{z} \geq 0$, then the unique equilibrium price vector is:*

$$\boldsymbol{p}^{*} = -\frac{1}{\phi}[\Omega + A(A'MA)^{+}A']^{-1}A(A'MA)^{+}A'\boldsymbol{\delta} \tag{16}$$

for $\Omega = \mathrm{diag}(\boldsymbol{a_1}'(A'MA)^{+}\boldsymbol{a_1}, \ \ldots, \ \boldsymbol{a_K}'(A'MA)^{+}\boldsymbol{a_K})$.

Proof: Consider a given firm i's objective function:

$$\pi_i(\boldsymbol{p}) = p_i q_i(\boldsymbol{p}) = p_i\big[\sum_j a_{ij}(A'MA)^{+}(\sum_k a_{kj}\delta_k + \phi a_{ij}p_i + \phi\sum_{k\neq i} a_{kj}p_k)\big] \tag{17}$$

The set of first-order conditions for this problem across all firms may therefore be defined as

$$0 = A(A'MA)^{+}A'\boldsymbol{\delta} + 2\phi\Omega\boldsymbol{p} + \phi(A(A'MA)^{+}A' - \Omega)\hat{\boldsymbol{p}} \tag{18}$$

To find the Bertrand (Nash) equilibrium, we look for the fixed point $\boldsymbol{p}^{*}$ such that $\boldsymbol{p} = \hat{\boldsymbol{p}} = \boldsymbol{p}^{*}$ - i.e. that point which defines the best response as a function of all competing firms' strategies. $\square$

Uniqueness is enabled by the following property:

[6] Under the usual smoothness assumptions, if the Jacobian of the demand function is negative semi–definite then the function is locally monotone non–increasing with respect to prices-that is, it satisfies the law of demand. For a positive semi-definite $A'(A'MA)^{+}A'$, $\nabla\boldsymbol{q} = \phi A(A'MA)^{+}A'$, which is negative semi-definite for $\phi < 0$.

Lemma 3: *$\Omega + A(A'MA)^{+}A'$ is invertible.*

What if all K differentiated products in the assortment are sold by a single multi-product monopolist retailer?

Proposition 4: *If a single multi-product firm sells all goods, the dimensionality constraint leads to a multiplicity of price equilibria expressed as*

$$\boldsymbol{p}^{*} = -\frac{1}{\phi}\big[W\big]^{+}A(A'MA)^{+}A'\boldsymbol{\delta} + (I - W^{+}W)\boldsymbol{y} \tag{19}$$

for

$$W = \Omega + A(A'MA)^{+}A' + A(A'MA)^{+}A' \circ G \tag{20}$$

where $\boldsymbol{y} \in \mathbb{R}^{K}$ is arbitrary and G is a symmetric hollow $K \times K$ matrix where each element $G_{ab} = 1$ if $a, b \in S$, 0 otherwise.

Proof: A multi-product firm determines a vector of prices for its goods by maximising the following profit function:

$$\pi_S(\boldsymbol{p}) = \sum_{i \in S} p_i[\sum_{j} a_{ij}(A'MA)^{+}(\sum_{k} a_{kj}\delta_k + \phi a_{ij}p_i + \phi \sum_{k \neq i} a_{kj}p_k)] \tag{21}$$

for S the subset of products $i = 1, \ldots, K$ sold by said firm. The FOC can be expressed in matrix form as

$$0 = A(A'MA)^{+}A'\boldsymbol{\delta} + \phi(2\Omega + (A(A'MA)^{+}A' - \Omega))\boldsymbol{p} + \phi(G \circ A(A'MA)^{+}A')\hat{\boldsymbol{p}} \tag{22}$$

To find the Bertrand (Nash) equilibrium, we operate in a manner identical to that pursued to find the equilibrium price in the single-product firms case. □

The inclusion of an ownership matrix is problematic for the identification of a unique price equilibrium, as we cannot say anything *ex-ante* about the definiteness of G. In fact, no closed-form solutions for multi-product oligopoly case exist for this model. However, for the monopoly case in particular, the multiplicity of price equilibria are not just payoff-equivalent, but *equilibrium-quantity equivalent*:

Proposition 5: *In the monopoly setting, the unique optimal demand, given equilibrium prices $\boldsymbol{p}^{*}$, is defined by the expression:*

$$\boldsymbol{q}^{*} = \frac{1}{2}A(A'MA)^{+}A'\boldsymbol{\delta} \tag{23}$$

while the aggregate profit and consumer surplus take on the following unique closed-form expressions:

$$\Pi = \sum_{i}^{N} \pi_i = \boldsymbol{q}'(\boldsymbol{p}^*)\boldsymbol{p}^* = -\frac{1}{4\phi}\boldsymbol{\delta}' A(A'MA)^+ A'\boldsymbol{\delta} \tag{24}$$

$$CS = \boldsymbol{q}'(\boldsymbol{p}^*)\boldsymbol{\delta} - \frac{1}{2}\boldsymbol{q}'(\boldsymbol{p}^*)M\boldsymbol{q}(\boldsymbol{p}^*) = \frac{3}{8}\boldsymbol{\delta}' A(A'MA)^+ A'\boldsymbol{\delta} \tag{25}$$

The proof is trivial once we note that $G_{monopolist} = A(A'MA)^+A' - \Omega$ for a single multi-product firm selling all goods besides the outside option, so $W_{monopolist} = 2A(A'MA)^+A'$. Expressions derived from this consumer choice model can be particularly useful to understand the impact on retail settings where a given product goes out of stock.

Lemma 4: *Following the stock-out of a good sold by a monopolist retailer, the resulting changes in optimal demand, aggregate profits, and consumer surplus have the following closed-form expression:*

$$\Delta\boldsymbol{q}^* = -\frac{1}{4}\begin{pmatrix} \sigma_{ii}\delta_i + \boldsymbol{\sigma}_i'\boldsymbol{\delta}_{-i} \\ \boldsymbol{\sigma}_i\delta_i \end{pmatrix} \tag{26}$$

$$\Delta\Pi = -\frac{1}{4\phi}\left[\sigma_{ii}\delta_i^2 + 2\delta_i(\boldsymbol{\sigma}_i'\boldsymbol{\delta}_{-i})\right] \tag{27}$$

$$\Delta CS = -\frac{3}{8}\left[\sigma_{ii}\delta_i^2 + 2\delta_i(\boldsymbol{\sigma}_i'\boldsymbol{\delta}_{-i})\right] \tag{28}$$

for $\Sigma = A(A'MA)^+A'$.

Even where M is symmetric and diagonally dominant (common assumptions, sufficient for positive definiteness), consumer surplus and aggregate profits may or may not *decrease* as a result of product exit, and the ways in which they do so will depend on the structure of the representative consumer's consideration set. The QQUM specification aids us in allowing the expressions above to be written in closed form. This may not be the whole story however; a product whose removal raises aggregate profit may nonetheless be kept within the assortment if it can minimise the impact of other products' removal. As can be seen above, the removal of any given product is exacerbated or dampened by its interaction with others. A correct understanding of the interaction matrix is important for addressing stock-outs.

6 Theoretical implications

6.1 Drivers

Key to the model above is the degree of linear dependency amongst the rows of A.

Lemma 5: *If A is full row rank, then* $A(A'MA)^+A' = M^{-1}$.

Just as how earlier linear dependence was sufficient, but not necessary, for the NN-LF condition to constrain $\boldsymbol{q}$, a full rank A is necessary, but not sufficient, for an unconstrained $\boldsymbol{q}$. This is because of the NN condition, $\boldsymbol{z} \geq 0$. When the full NN-LF condition binds,

$$\begin{aligned} \boldsymbol{q} &= A(A'MA)^{+}A'(\boldsymbol{\delta} + \phi\boldsymbol{p}) \ \ s.t. \ \ \boldsymbol{z} \geq 0 \\ &= A(A'MA)^{+}A'(\boldsymbol{\delta} + \phi\boldsymbol{p}) + A(A'MA)^{+}\boldsymbol{\lambda}(\boldsymbol{p}) \\ &= A \cdot \arg\min_{z \geq 0} ||Az - A(A'MA)^{+}A'(\boldsymbol{\delta} + \phi\boldsymbol{p})||_M^2 \end{aligned} \tag{29}$$

When only the NN portion binds, $span(A) = \mathbb{R}_+^K$, $A(A'MA)^{+}A' = M^{-1}$, and

$$\begin{aligned} \boldsymbol{q} &= M^{-1}(\boldsymbol{\delta} + \phi\boldsymbol{p}) \ \ s.t. \ \ \boldsymbol{z} \geq 0 \\ &= M^{-1}(\boldsymbol{\delta} + \phi\boldsymbol{p}) + A(A'MA)^{+}\boldsymbol{\lambda}(\boldsymbol{p}) \\ &= A \cdot \arg\min_{z \geq 0} ||Az - M^{-1}(\boldsymbol{\delta} + \phi\boldsymbol{p})||_M^2 \end{aligned} \tag{30}$$

In either case, $\lambda_i > 0$ exactly when the corresponding basket coordinate binds at the lower bound $z_i = 0$. The matrix $A(A'MA)^{+}A'$ (or M^{-1}, respectively) maps those shadow costs into product space; its image is a cone spanned by the columns of A associated with the binding coordinates. We may call this additional expression a price-contingent corner-solution adjustment. When $\boldsymbol{\lambda} = \boldsymbol{0}$, the initial closed-form solutions hold. Significantly, $\boldsymbol{\lambda}$ is a function of price; we return to this later to discuss its implications for how to address price endogeneity in an empirical setting.

We have yet to spell out the mechanism that is driving the difference in competitive behaviour under different consideration set structures. Firstly, note the following, once again under the assumption that the NN condition is satisfied:

Lemma 6: *The satisfaction of the LF condition alters the demand sensitivity to any price change $\Delta\boldsymbol{p}$ that is not in the column space of A.*

Can we say anything about how M^{-1} differs structurally from $A(A'MA)^{+}A'$? Yes:

Proposition 6: *Where the LF condition binds, relative to M^{-1}, $A(A'MA)^{+}A'$ has (weakly) lower-valued diagonal elements.*

Proof: The trick for this proof is to use congruence. Note that $M^{-1} = M^{-1/2}IM^{-1/2}$ and $A(A'MA)^{+}A' = M^{-1/2}PM^{-1/2}$, for $P = \tilde{A}(\tilde{A}'\tilde{A})^{+}\tilde{A}'$ where $\tilde{A} = M^{1/2}A$ and P is an orthogonal projection onto the column space of $\tilde{A}$. By the aforementioned properties of the M matrix, this means M^{-1} and I are congruent, and so are $A(A'MA)^{+}A'$ and P. For our purposes, this allows us to compare P and I knowing that said comparison extends to their congruents. For example, consider the diagonal elements of P. The definiteness properties of I and P's congruents extend to them: $\boldsymbol{x}'P\boldsymbol{x} \leq \boldsymbol{x}'I\boldsymbol{x}, \forall\boldsymbol{x}$. Since $\boldsymbol{x}'I\boldsymbol{x} = \boldsymbol{x}'\boldsymbol{x} = ||\boldsymbol{x}||^2$, $P_{ii} = \boldsymbol{e}_i'P\boldsymbol{e}_i \leq \boldsymbol{e}_i'\boldsymbol{e}_i = ||\boldsymbol{e}_i||^2 = 1$. As all diagonal elements of I are equal to 1, this means all diagonal elements of P are weakly smaller than I's, and all diagonal elements of $A(A'MA)^{+}A'$ are weakly smaller than M^{-1}'s (one need only think of $\boldsymbol{x} = M^{-1/2}\boldsymbol{e}_i, \forall i$). Each individual diagonal element of $A(A'MA)^{+}A'$ is weakly smaller than the diagonal elements of M^{-1}, so the average of these is also smaller. □

This attenuation effect is intuitive. Where consumers purchase combinations of goods, individual goods and their characteristics will drive their own demand less, as that will depend on the wider set being purchased. In these cases, own price makes less of an impact than it would otherwise.

This result can be extended also to the NN condition, which imposes an attenuation effect of its own.

Corollary 1: *Where the NN condition binds, own-price effects are (weakly) smaller in magnitude.*

The argument follows the Proposition 6 proof. The NN condition introduces an active set of baskets with $z_j > 0$ and separates it from an inactive set where $z_j = 0$. This effectively replaces the column space of A with a smaller (active) subspace in the column space of $A_{z_j>0} \subseteq A$. As we saw, orthogonal projection onto a smaller subspace weakens the length of any component's projection, thus leading to diagonal entries (weakly) smaller in magnitude.

Important questions remain. No stronger claim can be made regarding the exact change in the characteristics or magnitude of the off-diagonal elements of $\phi A(A'MA)^{+}A'$ relative to ϕM^{-1}. Similarly, not much can be said of the impact of the different characteristics of the model on its equilibrium outcomes.

6.2 Simulation

To address these limitations, I run a Monte Carlo simulation of product–basket data. This allows us to systematically generate random environments, solve for equilibrium quantities and prices under different market structures, and evaluate different consideration set structures. While not constituting solid proof, these outcomes may provide some degree of confidence as to the direction of some otherwise hard-to-prove effects, and constitute hypotheses for further research. The simulation is run over 1000 draws, and is divided into three steps: primitive sampling, interaction matrix formation, and equilibrium computations.

In the first step, the representative consumer's consideration set A is produced from a baseline of 10 goods and 51 baskets (roughly 5% of all feasible combinations). After correcting for zero- and non-unique columns and rows via non-negative linear combinations of the remaining, up to 10 columns and rows are also added in such a way. The average draw has 14.8 goods and 55.5 baskets, for an average rank of 10 and an average 14.3 missing standard basis vectors. The marginal-utility-of-income coefficient ϕ is set to -0.5.

In the second step, the consideration set A is leveraged to produce the interaction matrix M. To do so, I implement the approach in Tian et al. (2021) to produce a matrix of substitute goods based on co-purchase regularities. This I discuss in greater detail below. The approach allows both A and M to be indirectly inferred from the same underlying simulated preferences. The full interaction matrix itself is 50% said specification, 50% randomised, to reflect possible measurement error. Taste correlations are adjusted for directly in the Tian et al. (2021) measure, which is explained in greater detail below.

In the last step, I compute four different models, underpinned by two market structures: multi-good monopoly and single-good monopolistic competition; and two choice models: over shopping baskets and over goods, the latter of which acts as a baseline.

After each simulation draw, all models are compared across a range of characteristics: given the effective price elasticity of demand in each draw - the Jacobian of equilibrium demands with respect to prices at the computed optima depends on the NN condition - I determine own-price elasticity, cross-price elasticity, average magnitude of substitutes' and complements' cross-price elasticities, and the average number of complementary and substitute products per good; and outcomes: average demand, consumer surplus, and aggregate profits.

The simulation results are divided into three categories: findings which hold across all simulations, and therefore carry a degree of confidence to their generalisability; findings which hold across nearly all simulations, and therefore can be trusted in settings comparable to the one simulated; and findings which hold across a small majority of all draws, and are therefore perhaps most revealing for their ambiguity.

Table 1: Simulation of consumer choice counterfactuals

Metric	**Constrained**	**Unconstrained**	**Relation**	**% draws**
Avg own-price effect	−0.37	−0.53	< (abs.)	100
Avg cross-price effect	0.008	0.0185	<	100
Avg substitution effect	0.0336	0.0323	>	65.7
Avg complementary effect	−0.0497	−0.0198	> (abs.)	92.4
Avg number of substitutes	4.7	5.2	<	76.6
Avg number of complements	2.12	1.69	>	73.1
Firm consumer surplus	89.3	178.1	<	95.6
Monopoly consumer surplus	106.5	107.4	<	95.7
Firm aggregate profit	369.7	391.2	<	62.2
Monopoly aggregate profit	426.1	429.5	<	99.4
Firm aggregate demand	2.59	3.73	<	94.2
Monopoly aggregate demand	2.88	2.90	<	95.6

Notes: "Relation" compares the outcomes of consumer choice under consideration set constraints to the unconstrained setting value. Rows marked "(abs.)" compare absolute values. "% draws" is the percentage of simulation draws in which the stated relation holds. No metric has an equality relation between models over 15% of all draws.

Two results were found to hold in all draws, and are expressed as follows:

"Strong" Hypothesis 1: *The absolute value of the average own-price effect is strictly lower when choice is constrained by the representative consumer's consideration set.*

This is as predicted in Proposition 6, though slightly stricter. More novel is the following:

"Strong" Hypothesis 2: *The average cross-price effect is strictly lower when choice is constrained by the representative consumer's consideration set.*

This constitutes one of the more elusive proofs of interest to us. Overall, the attenuation effect found in own-price effects appears to extends to cross-price effects. A set of "weak" hypotheses offers some suggestions on the drivers of this result.

"Weak" Hypothesis 1: *The average magnitude of cross-price effects from complementary goods is tendentially greater when choice is constrained by the representative consumer's consideration set.*

The average magnitude of cross-price effects from substitute goods was found to be greater under constrained choice in a small majority of all simulation draws.

"Weak" Hypothesis 2: *The average number of complementary goods per product is generally greater when choice is constrained by the representative consumer's consideration set.*

"Weak" Hypothesis 3: *The average number of substitute goods per product is generally lower when choice is constrained by the representative consumer's consideration set.*

These suggest that the decrease in the magnitude of the cross-price effects under constrained choice is driven by a simultaneous decrease in the number of substitutes (and to some extent the strength of the respective cross-price effects) and increase in the number of complements and respective cross-price effects.

What may be going on here? At a base level, some goods effectively become more or less attractive, closer or farther substitutes, or even complementary, as cross-price effects become intertwined with often-co-purchased goods. Consider the earlier $A_{\{milk,bacon,pasta\}\times choices} = [2\ 0\ 1\ 2;\ 2\ 2\ 2\ 4;\ 0\ 2\ 1\ 2]$. The Tian et al. (2021) approach for producing a substitute matrix, combined with matrix $A_{\{milk,bacon,pasta\}\times choices}$, produces the following matrices:

$$M^{-1} = \begin{pmatrix} 1.140 & -0.107 & -0.370 \\ -0.107 & 1.030 & -0.107 \\ -0.370 & -0.107 & 1.140 \end{pmatrix},\ A(A'MA)^{+}A' = \begin{pmatrix} 0.884 & 0.257 & -0.627 \\ 0.257 & 0.514 & 0.257 \\ -0.627 & 0.257 & 0.884 \end{pmatrix} \quad (31)$$

Notice how, despite all three items behaving as substitutes in the first matrix, bacon, the weakest substitute, becomes a complement in the second matrix. Milk and pasta meanwhile become stronger substitutes. This seems to be driven by bacon being in all consideration set alternatives, so all feasible ways for the consumer to adjust what they buy run through it. Once choice is constrained, this fact makes bacon a complement to both milk and pasta; conditional on bacon, the relevant margin is what to pair it with, milk or pasta. This residual rivalry explains the strengthening of substitute effects. Effectively, regular co-consideration appears to drive complementarities, while strengthening substitution between goods which share co-considered products.

What do these results spell in terms of competitive outcomes?

"Weak" Hypothesis 4: *Consumer surplus is generally lower when choice is constrained by the representative consumer's consideration set, across both single-product oligopolistic and multi-product monopoly settings.*

Product ownership concentration was found to lead to greater consumer surplus where consideration set structure binds in a majority of simulation draws, a result virtually absent in product choice. This appears to be due to the above-mentioned complementarities created by co-consideration, best taken advantage of in concentrated markets where incentives are internalised.

"Weak" Hypothesis 5: *Aggregate profit is generally lower when choice is constrained by the representative consumer's consideration set, in the multi-product monopoly setting.*

Surprisingly, however, single-product firms yield greater aggregate profits under constrained choice in a small majority of cases. This may be driven by the two "strong" hypotheses above: weaker own-price effects result in higher pricing power and thus higher prices; smaller cross-price effects - which as we saw are driven by a greater (lesser) number of complements (substitutes) and complementary (substitution) effects - diminish competitive concerns. The strategic complementary characteristic of Bertrand competition does the rest. Nonetheless, aggregate profits are always greater under multi-product monopoly than in single-product oligopolistic settings.

"Weak" Hypothesis 6: *Average demand is generally lower when choice is constrained by the representative consumer's consideration set, across both single-product oligopolistic and multi-product monopoly settings.*

The rational inattention literature consistently finds that behavioural restrictions to consumer choice reduce price elasticity (Matějka and MacKay, 2015; Caplin, Dean, and Leahy, 2019). This is intuitive for our setting too. Where choices are restricted by the NN-LF, for every consideration set alternative its next best alternative is on average worse than if there was no such restriction on the consumption bundles. Unlike in the existing literature, this effect is not necessarily driven by an "information gap", but can instead be a consequence of the underlying structure of preferences itself. Constrained choice takes on an appearance of making the average good less attractive for consumers.

Nonetheless, average demand under constrained choice is greater under a multi-product monopoly than in single-product oligopolistic settings in a small majority of draws. This is never the case otherwise.

Of the above, 96 simulation draws contain a full-rank A. This allows for an additional discussion of the case where only the NN condition binds.

Both the "strong" hypotheses hold also under the NN condition alone, as does "weak" hypothesis 1. The number of complements and substitutes remains unchanged. Aggregate profits, average demand, and consumer surplus all fall, though no difference is recorded between the different competitive settings. However, the effects found are minimal across the board. Most likely, this is because the wedge created by the corner-solution adjustment is only a problem for a researcher estimating price effects, leaving the attenuation effect - without a corresponding scramble of interaction effects driven by the LF condition - as the sole impact. A partial relationship between A and M also minimises the degree to which the NN condition binds.

On the whole, when the structure of the consideration set binds, own-price effects are weaker than they would otherwise be, as goods' demand becomes relatively more dependent on those goods with which they are often co-purchased. Often co-considered goods behave as complements, while goods often paired with the same products face substitution effects conditional on these. From these, consumer choice under more realistic consideration set constraints also translates into worse market outcomes for both consumers and firms than previous theory otherwise assumed; and it favours product ownership concentration at a greater clip, with the average draw repeatedly finding greater aggregate demand, consumer surplus, and aggregate profits in this setting.

Most of the findings that have been presented relate to the LF condition. The condition will always bite if consumers consider more products across shopping baskets than the total number of shopping basket alternatives considered. Outside of this setting, how often is A less than full row rank in reality? Even if linear dependence is the case for a single individual's consideration set, aggregation into a representative consumer will likely change this. Even if individual purchase patterns are driven by latent factors that create linear dependency, aggregation may bring A to full rank. Each individual might even be influenced by a similar set of factors, but should the factors hold different weights depending on the consumer, this heterogeneity may blur the low–rank structure. An aggregated A matrix with contributions from many individuals with differing factor loadings will have increasing effective diversity (or rank). In other words, while each individual's consideration set might be low rank, averaging over many individuals could add enough independent variation that the aggregate no longer possesses said dependencies.

To explore the impact of the NN-LF conditions on real aggregated data, I retain the example of the supermarket industry, and proceed to consider an empirical application. In doing so, I make use of the capabilities of linear demand and demonstrate how such a model may be estimated over a large product assortment, while minimising potential concerns raised by previous research.

7 Empirical application

7.1 Data

In keeping with the twin goals of tractability and parsimony made possible by the QQUM specification, I make use of detailed point-of-sale data. The data was obtained from *SONAE MC*, a leading supermarket chain in the Portuguese food retail market. The dataset represents roughly 5% of all transactions conducted across four supermarkets between February 2020 to February 2023. It covers date and location of each transaction; the set of products purchased per transaction both at the product level and across 4 different product category layers; and units purchased, gross sales value, and discounts issued per product per transaction. All data is collected through the stores' scanners in the moment of the sale.

The four supermarkets - Supermarket Maia, Hypermarket Mafra, Hypermarket Évora, and Hypermarket Portimão - were selected prior to the data analyses out of a list of 353 shops given their geographical distribution, distance from competitors, and size heterogeneity. The selection along these lines was made to match the model's assumptions - the primacy of intra-supermarket competition in particular - while including sufficient heterogeneity to consider the internal validity of the model's outcomes under different conditions. An evaluation of supermarket participation rates - i.e. the transaction volumes observed each period per supermarket - suggests that, for our sample of grocery stores, inter-supermarket competition is not a major concern (details are in the Appendix).

Table 2: Transaction statistics across products and categories

	Products	**Categories**
A: Transaction size (number of items)		
Mean	60.43	5.12
Standard deviation	67.99	3.69
B: Percentile transaction size		
25th	16	2
Median	36	4
75th	78	7
C: Transaction share by size class		
$\leq$ 20 unique items / $\leq$ 2 groups	32.97%	31.04%
21-80 unique items / 3-8 groups	42.80%	50.32%
81-140 unique items / 9-14 groups	13.74%	16.97%
$>$140 unique items / $>$14 groups	10.49%	1.67%
D: Revenue share by transaction size class		
$\leq$ 20 unique items / $\leq$ 2 groups	8.04%	8.17%
21-80 unique items / 3-8 groups	30.52%	40.91%
81-140 unique items / 9-14 groups	23.21%	42.56%
$>$140 unique items / $>$14 groups	38.23%	8.36%

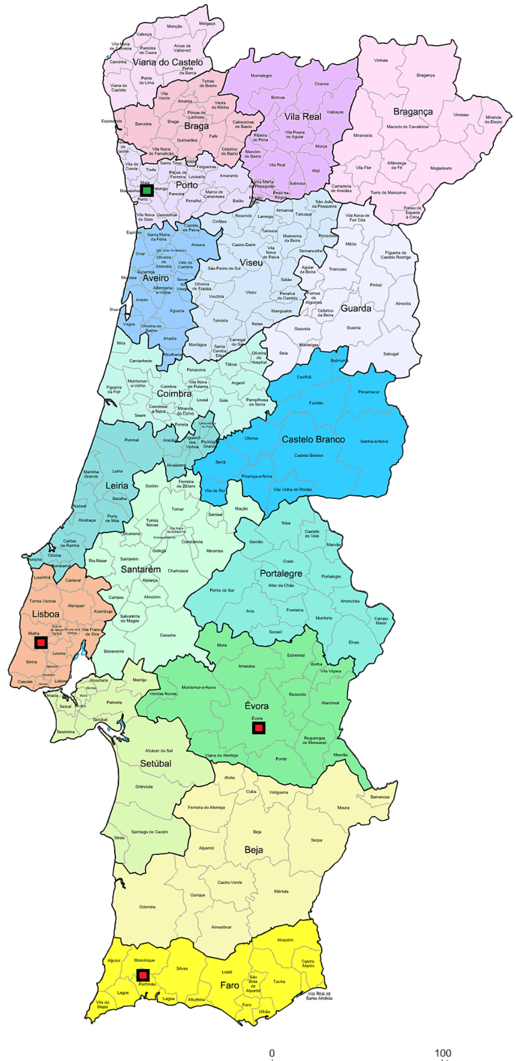

Figure 1: Shop locations - hypermarket (red), supermarket (green). Hypermarkets are major retail hubs, meant to attract consumers to a distinct location within the urban fabric. Supermarkets are mid-sized stores targeting proximity to consumers.

Table 3: Shop statistics

Shops by Type	Mean transaction size	SD transaction size
A: Hypermarket		
Évora	70.85	77.90
Mafra	52.68	58.12
Portimão	60.52	68.50
B: Supermarket		
Maia	49.00	51.35

Table 4: Product availability in stores and the average revenue brought by products at each level of availability

	Products	Avg Revenue	Categories	Avg Revenue
At only 1 store	2%	€ 2 206.4	5%	€ 1 334.0
At 2 stores	3%	€ 693.5	-	-
At 3 stores	19%	€ 487.3	8%	€ 1 308.0
At all stores	76%	€ 970.8	86%	€ 644 261.5

Table 5: Shop statistics - note that all stores share 18910 products

	#	% Unique	% Revenue	% Unique Revenue
A: Hypermarket				
Évora	24 069	0.41%	31.91%	0.56%
Mafra	23 896	0.11%	13.73%	0.04%
Portimão	24 564	0.97%	44.00%	0.44%
B: Supermarket				
Maia	19 442	0.33%	10.36%	0.11%

The following procedures were applied to filter outliers and products for which insufficient data was made available. Any products with (1) fewer than 100 transactions, (2) sold at zero price, and/or (3) without any price variation were removed. Furthermore, for each product, observations pertaining to a given month where no close competitor was found across supermarkets other than that product's are also discarded. Lastly, any products which go unobserved across a given time period in a given supermarket are assumed to not have been available for purchase.

Table 6: Product category statistics

Categories	Products	Store presence	Sales share
Groceries (Salty)	1 766	4	4.94%
Groceries (Sweet)	2 228	4	4.64%
Soft Drinks	1 267	4	7.09%
Crates	9	4	0.01%
Hygiene	1 597	4	4.21%
Home Cleaning	2 173	4	7.18%
Frozen Goods	758	4	3.27%
Dairy and Derivates	1 478	4	7.90%
Beauty	1 665	4	2.94%
Essential Goods	690	4	3.24%
Butchery	452	4	6.35%
Fishmonger	558	4	6.47%
Charcuterie and Cheese	1 182	4	6.85%
Breakfast	1 564	4	5.55%
Fruits and Vegetables	1 037	4	10.29%
Bakery	999	4	4.35%
Wine and Spirits	1 019	4	4.83%
Take Away	468	4	2.30%
Luggage and Sport	1	4	0.00%
Bio and Healthy	6	4	0.00%
Culture	884	4	0.96%
DIY	744	4	1.14%
Petfood and Care	798	4	2.15%
Home Goods	1 250	4	1.75%
Stationery/Newsagents	80	4	0.18%
Women Apparel	1	3	0.00%
Adult Non-Apparel	2	1	0.00%
Nursery	13	4	0.01%
Large Domestic Goods	1	3	0.01%
Small Domestic Goods	12	4	0.03%
Others	223	4	1.32%
TOTAL	24 925		

This correction leaves us with 24 925 goods remaining, spanning 37 major categories, 118 sub-categories, and 571 sub-sub-categories. Each product is observed on average 1 210 times over the sample, with 175, 369, and 937 observations at the 25th, median, and 75th percentiles, respectively.

We observe 456 439 unique transactions - defined by their goods down to the number of units per product - out of 499 047 baskets in total. The average number of observations per unique basket is 1.093; 2.16% of unique baskets are observed more than once, accounting for 10.51% of all baskets. Consistent with the shop-level results, hypermarkets exhibit larger and more dispersed basket sizes, stock more staple and unique items, and command a disproportionate revenue share. Aggregate gross sales equal € 82 636 508. Discounts total € 156 601, representing 0.19% of gross sales.

It was mentioned earlier that the NN-LF conditions holds if and only if not all goods are considered in standalone consideration set alternatives. In this sample, 45.8% of all goods are not purchased standalone. In what follows, I assume that this matches the structure of the true, latent, consideration set of the set of shoppers in the sample. More on this later.

7.2 Identification Strategy

I estimate the set of linear demand equations micro-founded from the structural model above, one per good with supermarket and time FEs. My approach follows Pinkse and Slade (2004) with some significant changes.

For a given market m - defined at the supermarket-quarter level we have, for the unconstrained consumer choice model:

$$\boldsymbol{q}_m = M_m^{-1}(X_m\boldsymbol{\beta} + \phi\boldsymbol{p}_m) + \boldsymbol{u} \tag{32}$$

for $\boldsymbol{\delta} = X\boldsymbol{\beta}$ (we consider latent product characteristics shortly). Matrix M is unknown.

I simplify the expression further by defining a series-expansion approximation of the inverse matrix, helpful both for model fitting but also to minimise numerical instability and computation concerns stemming from inversion.

$$M^{-1} \approx g(W) = \sum_{j=0}^{J} \alpha_j W^j \tag{33}$$

for W a proxy for cross-price effects, and J the number of series expansion terms. More on proxy selection later. Under this specification, we may write the model as follows:

$$\begin{aligned}\boldsymbol{q}_m &\approx \sum_{j=0}^{J} \alpha_j W_m^j (X_m\boldsymbol{\beta} + \phi\boldsymbol{p}_m) + \boldsymbol{u} \\ &= X_m\boldsymbol{\beta} + \phi\boldsymbol{p}_m + \sum_{j=1}^{J} \alpha_j W_m^j (X_m\boldsymbol{\beta} + \phi\boldsymbol{p}_m) + \boldsymbol{u}\end{aligned} \tag{34}$$

The formulation in (34) shows how the original regressors $X\boldsymbol{\beta}$ and $\phi\boldsymbol{p}$ are augmented by so-called "spatial" lags $W^j(X_m\boldsymbol{\beta} + \phi\boldsymbol{p}_m)$ weighted by the coefficients α_j. In practice, this means that our estimation model can be set up as a standard linear regression that includes both the original variables and their corresponding "spatially" lagged versions. This is computationally favourable as it allows "spatial" lags to be precomputed, and is effectively a series of linear combinations, which we may estimate via standard techniques. It also enables the separate identification of α_j from ϕ if so desired, via the delta method.

Stacking across supermarkets, and including covariate interactions with prices to account for coefficient heterogeneity, the final model specification is as follows:

$$\boldsymbol{q} \approx X\boldsymbol{\beta} + (\overline{\phi} + X\boldsymbol{\eta})\boldsymbol{p} + \sum_{i=1}^{M} \mathbf{1}\{i = m\} \left[\sum_{j=1}^{J} \alpha_j W_m^j (X\boldsymbol{\beta} + (\overline{\phi} + X\boldsymbol{\eta})\boldsymbol{p}) \right] + \boldsymbol{u} \qquad (35)$$

Term $\boldsymbol{u}$ is a vector of regression errors, cluster-robust at the supermarket-quarter level.

How does consideration set structure impact linear demand estimation? Consider now the implications of the NN-LF constraint to this empirical strategy. The best-case scenario is that where only the NN portion of the constraint binds. From (30), $\boldsymbol{z} \geq 0$ enters the Lagrangian via the multiplier λ, leading to a restated $\boldsymbol{q}^\dagger(\boldsymbol{p}, \boldsymbol{\lambda})$:

$$\boldsymbol{q}^\dagger = M^{-1}(\delta + \phi \cdot \boldsymbol{p}) + A(A'MA)^+\lambda(\boldsymbol{p}) \qquad (36)$$

If $\boldsymbol{\lambda} \geq 0$ and binds, it creates a wedge between $\boldsymbol{q}^\dagger$ and our original model $\boldsymbol{q}$: $\Delta\boldsymbol{q} = \boldsymbol{q}^\dagger - \boldsymbol{q} = A(A'MA)^+\boldsymbol{\lambda}$. This wedge is our earlier price-contingent corner-solution adjustment. For demand estimation, this is problematic. A price-dependent $\boldsymbol{\lambda}$ introduces a new form of endogeneity hitherto disregarded. Because $\boldsymbol{\lambda}$ is unobserved, we cannot instrument said endogeneity away.

If the NN condition binds, how can we de-endogenise prices? First, I compute $\hat{\boldsymbol{q}}^\dagger$, the fitted values of an initial "naive" regression and estimate $\hat{M^{-1}}$. Second, I find $\boldsymbol{r} = \hat{\boldsymbol{q}}^\dagger - A \cdot \arg\min_{\boldsymbol{\lambda} \geq 0} ||A\boldsymbol{\lambda} - \hat{\boldsymbol{q}}^\dagger||^2_{\hat{M}}$ via iterated weighted Non-Negative Least Squares (NNLS) estimation. Third, after I get $\boldsymbol{r}$, I plug the result back onto the LHS of the regression so the dependent variable becomes the observed $\boldsymbol{q} + \boldsymbol{r}$ (or $\boldsymbol{y} + \boldsymbol{r}/\mathrm{sd}_Q$ in standardised units), i.e., a control-function step. Lastly, I re-run the regression. This removes the price-dependent corner-solution component from the error term while avoiding new endogeneity from $\boldsymbol{\lambda}$ because the correction residual $\boldsymbol{r}$ is $\hat{M}$-orthogonal to the span of A.

What if the full NN-LF condition binds? We have assumed thus far that, for an appropriate proxy W, we have that $M^{-1} \approx \sum_{j=0}^{J} \alpha_j W^j$. What is less clear is whether W is an appropriate proxy for M^{-1} or if it should be used to approximate $A(A'MA)^+A'$ as a whole where the NN-LF binds. This is a matter better left to a test: model both cases and verify which fits the data best. This is less so a problem in the latter case, as our estimation approach above already fits that alternative without any additional changes. The former, however, requires expressing $A(A'MA)^+A'$ as a function of M^{-1}, so that $\sum_{j=0}^{J} \alpha_j W^j$ can be fit.

To achieve this, note:

Lemma 7: *For M a positive definite matrix, $A(A'MA)^+A' = M^{-1}A(A'M^{-1}A)^+A'M^{-1}$.*

If $A(A'MA)^+A' = M^{-1}A(A'M^{-1}A)^+A'M^{-1}$, then we can implement grid search to jointly estimate the parameters in $M^{-1} \approx \sum_{j=0}^{J} \alpha_j W^j$: set initial values for $\alpha_j, \forall j$, calculate $\hat{M^{-1}}A(A'\hat{M^{-1}}A)^+A'\hat{M^{-1}}$, and regress the model to obtain $\hat{\boldsymbol{\beta}}$ and $\hat{\phi}$. The optimal $\boldsymbol{\alpha}$

will be that which best fits the data.

To proxy for M, I take advantage of the frequency with which products are considered for joint purchase. In the economics literature, this is an approach closest to that in Atalay et al. (2023) and with roots in stochastic choice (e.g. Manzini, Mariotti, and Ülkü, 2019). The proxy itself is drawn from Tian et al. (2021), who approach the matter of cross-price effect estimation from a network science perspective. The authors use point-of-sale data and estimate complementarity and substitution relationships by the frequency with which goods are jointly purchased, relative to what would otherwise be expected.

A modified approach to that used in Tian et al. (2021) follows.[7] Complements are defined as goods jointly purchased more often than expected, with the degree correlated to the extent of this. Substitutes are goods that both share the same complements and are jointly purchased less often than expected, with the degree measured similarly. The sales data is converted into a product-transactions matrix S with every (row) good matched to every (column) transaction it appears in. Note that "transactions" by Tian et al. (2021) refers to non-unique "baskets". Earlier we have expressed the size of the set of all unique baskets as J; let us set T as the total number of observed transactions. From matrix S, I compute two similarity measures.

The first similarity matrix – equivalent to the *original measure* in Tian et al. (2021) - represents how similar the purchase patterns are between two products, yielding a measure of complementarity:

$$W^{(c)} = A^{(c)} \circ \cos(\theta_c) \tag{37}$$

The term $\cos(\theta_c)$ is a matrix whose elements are defined as, for a pair of goods $a, b \in \mathcal{K}$:

$$\cos(\theta_c)_{ab} = \frac{\Xi_{ab}}{\sqrt{\Xi_{aa}}\sqrt{\Xi_{bb}}} \tag{38}$$

for $\Xi = (D^{\mathcal{P}})^{-1} S (D^{\mathcal{T}})^{-1} S' (D^{\mathcal{P}})^{-1}$, $D^{\mathcal{P}}$ a diagonal matrix of the product degrees (how many transactions per good), and $D^{\mathcal{T}}$ a diagonal matrix of the transaction degrees (how many goods per transaction) . $D^{\mathcal{P}}$ normalises the product vectors, whereas $D^{\mathcal{T}}$ weighs the relative importance of each transaction as the inverse of transaction size - smaller shopping baskets in the consideration set are assumed to pertain to stronger ties between the included goods. The matrix Ξ is therefore a weighted version of the product-transactions matrix S.

[7]To validate their results, the authors utilise a dataset of flavour compounds in ingredients and a dataset of recipes to compute the relative number of shared flavour compounds between pairs of products. They find that pairs assessed to be substitutes have a significantly higher probability of sharing all their flavour compounds and a lower probability of appearing in the same recipes compared to all product pairs. Conversely, they observe that complementary pairs have a significantly higher probability of sharing no flavour compounds and a higher probability of appearing in more of the same recipes. Joint recipe presence and shared flavour compounds are found to yield approximately the same complement and substitute role assignments as those derived from the sales data.

The matrix $\cos(\theta_c)$ is defined under the assumption that all goods are complementary. It must therefore be supplemented by an analysis of joint purchases to determine the degree to which their frequency is *significantly* different from what would be expected of independent goods. To achieve this, we produce $A^{(c)}$, the matrix of potential complements, with elements equal to one in case of significance.

Significance is assessed under a statistical null model - the Bipartite Configuration Model (BiCM) in this case - whose goal is to create a randomised product-transactions matrix that preserves the observed degree sequence (i.e. the number of connections) for both products and transactions, while randomising all other structural properties. The BiCM is favoured in our setting as it takes the bipartite feature of the product-transactions matrix into account.

In the literature dedicated to incorporating combinations of goods (not just individual goods) as options within a discrete-choice setting, the key concern lies in separating co-purchases which result from preferences vs taste correlation. Correlated purchases can be the outcome of things like minimising trip-related costs or time and stock considerations. Preferences toward co-purchases meanwhile may come from a love for variety or complementarities derived from various setting mentioned above. Both the equation for $\cos(\theta_c)$ and the BiCM are meant to adjust for product popularity and transaction size precisely to minimise the influence of correlations.

For each product pair a and b the expected number of common transactions μ_{ab} is computed as:

$$\mu_{ab} = \frac{d_a^{\mathcal{P}} d_b^{\mathcal{P}} (\frac{1}{T}\sum_{t=1}^{T}(d_t^{\mathcal{T}})^2 - \frac{1}{T}\sum_{t=1}^{T} d_t^{\mathcal{T}})}{\frac{1}{K}(\sum_{t=1}^{T} d_t^{\mathcal{P}})^2} \tag{39}$$

for T the number of transactions, and $d_i^{\mathcal{P}}$ the i-th diagonal element of $D^{\mathcal{P}}$. Once we have μ we can compute the observed number of joint purchases $C = SS'$. Lastly, we compare c_{ab} with μ_{ab} using the Poisson distribution, on the assumption that in a sparse setting the probability of any two products co-occurring in a transaction is small:

$$A_{ab}^{(c)} = \begin{cases} 1 \text{ if } 1 - F_{ab}(c_{ab} - 1) < \alpha_c \\ 0 \text{ otherwise} \end{cases} \tag{40}$$

where $F_{ab}(y)$ is the Poisson CDF with mean μ_{ab} and α_c is a chosen significance threshold.

The second similarity matrix – equivalent to the *original substitutability measure* in Tian et al. (2021) - represents how similar pairs of products' complements are, yielding a measure of substitutability:

$$W^{(s)} = I\{(A^{(c)})'A^{(c)} > 0\} \circ A^{(l)} \circ \cos(\theta_s) \tag{41}$$

The expression $\cos(\theta_s)$ can be straightforwardly derived from $\cos(\theta_c)$. For every given product pair:

$$\cos(\theta_s)_{ab} = \frac{(\cos(\theta_c)\cos(\theta_c)')_{ab}}{\sqrt{\sum_{k=1}^{K}\cos(\theta_c)_{ak}}\sqrt{\sum_{k=1}^{K}\cos(\theta_c)_{bk}}} \tag{42}$$

The matrix $A^{(l)}$ behaves similarly to $A^{(c)}$. A pair of products is flagged as having significantly less joint purchases than would be otherwise expected of independent goods if:

$$F_{ab}(c_{ab}) < \alpha_l \tag{43}$$

where $F_{ab}(y)$ is once again the Poisson CDF with mean μ_{ab} and α_l is a chosen significance threshold. Should this be the case, the expression $I\{(A^{(c)})'A^{(c)} > 0\}$ is meant to identify the presence of shared complementary goods, and $\cos(\theta_s)$ provides the weights.

In both cases, the diagonal of each proxy (and their lags) is hollowed out to centre own-price effects in the second term of (34). The two similarity measures are tested both simultaneously and individually in the regression. The one providing the best fit turns out to be Tian et al. (2021)'s substitutability measure $W^{(s)}$, the results for which are provided below.

One of the crucial assumptions of linear regression when performing standard estimation is that the regression error $\boldsymbol{u}$ is exogenous i.e. orthogonal to, which is to say uncorrelated with, the set of explanatory variables X: $E(X\boldsymbol{u}) = 0$. When this is not the case, the parameters associated with those explanatory variables are neither consistent nor unbiased, respectively. One common concern in demand estimation is that of price endogeneity due to simultaneity bias - i.e. $E(\boldsymbol{p'u}) \neq 0$. Prices determine demand by consumers, but demand determines the prices set by firms. A naive demand estimation procedure could result in endogeneity.

Another common source of price endogeneity is omitted characteristics in product-level data. So far, we have remained agnostic on this. However, if $\boldsymbol{\delta} = X\boldsymbol{\beta} + \boldsymbol{v}$, the OLS condition becomes $E(\boldsymbol{p}'(M^{-1}\boldsymbol{v} + \boldsymbol{u})) = 0$. Since it is to be expected that prices and product characteristics are choices made simultaneously by sellers during the production stage and ahead of a sale, even without M^{-1}, $\boldsymbol{p}$ and $\boldsymbol{v}$ may be correlated. In this case, if one wishes to still use a least-squares estimator unchanged, either one assumes that the effect of those omitted characteristics is mean zero and independent across agents, or that said characteristics are orthogonal to (i.e. linearly independent from) the included characteristics. The latter assumption is rather tough to make in an environment where likely all product characteristics were chosen (close to) simultaneously but the first is rather commonplace. In such cases, sampling error tends to be unusually large, suggesting a model with low external validity and overfitting at best.

As such, both the simultaneity and the omitted characteristics problem is usually tackled via instruments for price, which may isolate exogenous variation in the endogenous variable and therefore determine our unobserved parameters. Berry, Levinsohn, and Pakes (1995) suggest using characteristics or cost shifters of either products themselves or their competitors. Noting these are often not available, Hausman (1996) proposes using prices of the same product in geographically distinct supermarkets. In the same vein, Nevo (2001) also explores price heterogeneity, as well as market characteristics, such as geographic heterogeneity more broadly defined (e.g. local labour markets). Blundell and Bond (1998) extend this to panel data, offering the option of lagged characteristics as less impacted by potential endogeneity concerns with national pricing, centralised inventories, or chain-wide discounts and promotions. Lastly, Fan (2013) instruments on aggregate consumer or market attributes, such as education or average income, amongst others.

I take advantage of the stream of observations on product purchases across five geographically distinct supermarkets each quarter of each year in our database, as well as the repeated cross-sectional nature of our data and our "spatial" lags, to define matrix Z of price instruments. I compute the average price of a given product's competitors across other supermarkets as our primary price instrument. Its lag, as well as the (lagged) number of competitors available within three increasingly narrowly-defined product categories, are also considered. These are potentially informative instruments as one may expect prices to be decreasing in the number of similarly categorised products available, and be affected by local economic dynamics. As exogenous variables, also in X, I include intercepts; supermarket and quarter FEs; private-label brand FEs; product category FEs at three levels of detail; and the number of supermarkets in our sample that each good is present in every time period. Lastly, following Pinkse and Slade (2004), for every additional "spatial" lag included in the model, additional instruments $W^j Z$ are computed.

The above is common practice, but in our setting the propriety of this approach will depend on the restrictiveness of an $E(M^{-1}\boldsymbol{v}|Z) = 0$ assumption. Insofar as $\boldsymbol{v}$ is composed of short-term product utility shocks, it is unlikely to play a significant role in defining long-term own- and cross-price effects in M^{-1} - i.e. the assumption is then equivalent to the usual exogeneity condition on IVs. This would not be the case if it were instead composed of latent product characteristics relevant to product competition. For the latter case, the rest of the paper operates under the assumption that such characteristics did not shift in a competitively significant way over our period of analysis. While it is too soon to tell, early indicators seem favourable (see, e.g. Lee, 2024).

All variables are standardised to ensure comparability and numerical stability.[8] A principal components analysis (PCA) is applied to X to orthogonalise covariates (minimising collinearity concerns) and extract the key factors driving variation in the data. Covariates are restricted to the smallest combination of components which contains 90% of X variation or higher.

The estimation procedure is carried out via Two-Stage Least Squares. All regressors, including price, are projected onto the instrument space. This is done by forming a matrix that includes both the direct regressors and their "spatial" lags, along with the instrument matrix built from the "spatial" lags. The projection step effectively purges the endogenous

[8] In the case of "spatial" lags, standardisation is performed after multiplication with the proxy, but the same moments are used.

variation in the regressors by isolating the variation that is explained by the instruments. The predicted values obtained from the first stage (i.e., the instrumented regressors) are then used in a pooled regression to estimate the structural parameters of interest. This stage yields consistent estimates by mitigating the bias that would arise from endogeneity.

After obtaining the parameter estimates and robust standard errors, I then compute t-statistics and p-values to assess the significance of the estimated coefficients.

7.3 Results

The revenue-weighted price index over the sample period is provided in Figure 2 for context, both in nominal and CPI-adjusted real terms.

Figure 2: Revenue-weighted mean prices

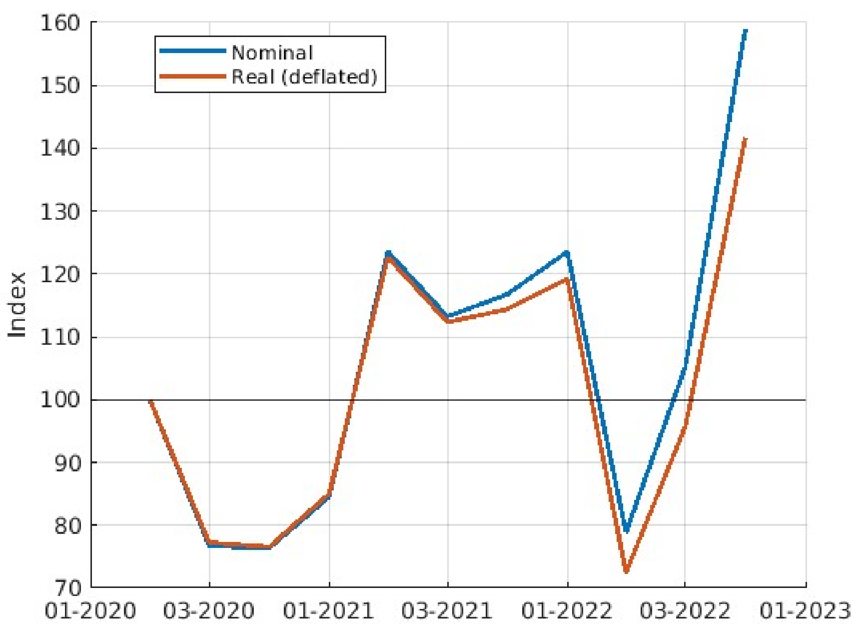

The substantial volatility observed tracks the timings of the four COVID waves in Portugal, with the price peaks generally following a given wave: the first starting in Q2 2020; the second in Q4 2020; the third in Q4 2021, and the fourth in Q2 2022.[9] The observed prices broadly follow the CPI, with a decoupling towards the end of the sample likely driven by energy prices and supply chain disruption, with global container freight indices - a benchmark for the average cost of shipping containers across major world trade routes - peaking at the

[9]This information was obtained from the WHO COVID dashboard at: https://data.who.int/dashboards/covid19/deaths?m49=620&n=o.

tail-end of 2021[10] and crude oil prices peaking during 2022.[11]

The results of our empirical strategy are presented below. The constraints due to consumer choice over shopping baskets bind via the NN constraint, while the LF constraint remains slack, facilitating estimation considerably.

Table 7: Model diagnostics and PCA summary

	Base model		NN-LF-adjusted model	
PCA				
# PCs for 90% variance	76		76	
Variance explained by top 10 PCs (%)	76.6		76.6	
Fit & diagnostics				
Adjusted R^2	0.63		0.12	
Wald F (122, 557514)	7 793.13	***	639.36	***
Durbin-Wu-Hausman χ^2 (33)	4 469.27	***	1 273.43	***

Notes: Robust standard errors in parentheses. *** : $p < 0.01$, ** : $p < 0.05$, * : $p < 0.1$. PCA is computed on the initial covariate set X. Degrees of freedom for the Wald test are shown as $(\mathrm{df1}, \mathrm{df2})$; the Durbin–Wu–Hausman statistic displays the χ^2 degrees of freedom in parentheses.

The PCA applied to X to orthogonalise covariates and extract the key factors driving variation in the data reduces the number of independent variables to 76. The top 10 principal components are used to form heterogeneous price coefficients in both specifications, and are shown to explain up to 76.6% of the variation across all covariates. Combined with time and place FEs as well as the "spatial" lags, the total number of regressors is 122, over 557,514 observations. The model retains substantial explanatory power even in its more restrictive version. The difference in fit reveals how much of the unadjusted results may be spurious. In each case, the included regressors are jointly informative, and the Hausman instrument - supported by its "spatial" lags - successfully minimises existing price endogeneity.

To analyse the implications of the NN-LF-adjusted model where it comes to price elasticities, an exhaustive analysis of the resulting mark-ups is presented below. The revenue-weighted mean mark-up is provided in Figure 3.

[10] See e.g. Monthly composite China Ningbo Container Freight Index at: https://www.statista.com/statistics/1320123/monthly-composite-china-ningbo-container-freight-index/.

[11] This information was obtained from Statista at: https://www.statista.com/statistics/326017/weekly-crude-oil-prices/.

Figure 3: Revenue-weighted mean mark-ups

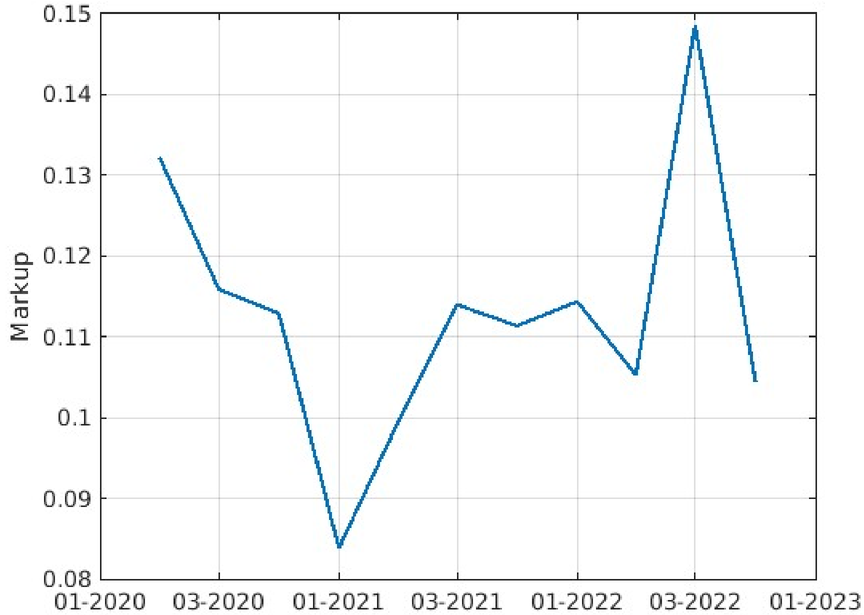

Mark-ups are calculated under the assumption that prices are set for each good independently, as the data does not specify ownership. The average mark-up hovers around 11% for most of the sample period, with a trough in the first quarter of 2021 (≈8.5%) and peak in the third quarter of 2022 (≈15%). These results correspond to own-price elasticities which match in magnitude those found elsewhere in the literature. In Dopper et al. (2024), a mean of median category-level mark-ups - which overweights less revenue-relevant, lower-mark-up goods relative to the measure presented here, and would thus supposedly be lower than our figure - points to an average mark-up of 63% in U.S. grocery store scanner data. The magnitude difference appears to be primarily driven by the difference in assumption of how products are priced.

Existing discrete-choice methods ignore cross-category effects. These have predictable implications for price elasticity estimates. Consider a stylised regression parameter for own-price effects, in the case where cross-category effects are ignored (or more broadly where market boundaries are understated):

$$\hat{\beta}_i = \frac{Cov(q_i, p_i)}{Var(p_i)} = \frac{\sum_k J_{ik} Cov(p_k, p_i)}{Var(p_i)} = J_{ii} + \sum_{k \neq i} J_{ik} \frac{Cov(p_k, p_i)}{Var(p_i)} \tag{44}$$

for $J = \phi \cdot A(A'MA)^{+}A'$. $J_{ii} < 0$ while J_{ki} may be positive or negative depending on the product pair. Firms are strategic complements under Bertrand competition, so retailer prices co-move weakly positively across substitute goods and weakly negatively otherwise, suggesting $Cov(q_i, p_i) \geq 0$ and $Cov(q_i, p_i) \leq 0$ respectively. The cross-price terms J_{ki} behave similarly, meaning that the summation contributes positively to $\hat{\beta}_i$. Thus, if cross-category price effects are ignored, own-price effect estimates are biased towards zero; this underestimates price elasticities, possibly producing excessive estimates of market power and understating pass-throughs.

The results obtained in the present paper match survey estimates of United States grocery store margins remarkably well.[12] U.S. data is used here due to the unavailability of quarterly figures for Portugal.[13] Figure 4 provides a side-by-side comparison of the four-quarter moving average mark-up in our sample and in the U.S. for our sample period.

Figure 4: Revenue-weighted four-quarter moving average mark-ups: sample & U.S. survey estimate

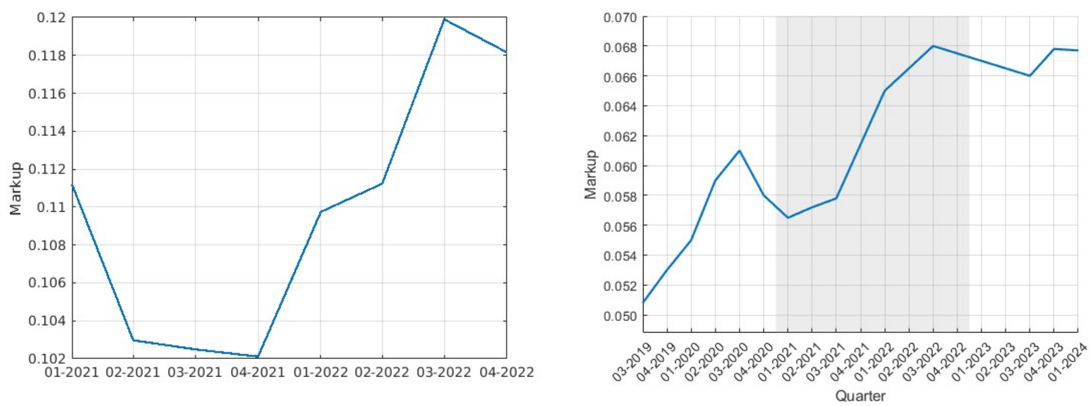

Notes: The grey band on the second figure identifies the time period of the first graph and our sample.

The trends are near identical, though levels differ cross-country; U.S. mark-ups see a trough of 5.7% in the first quarter of 2021 and a peak of 6.8% in the third quarter of 2022, roughly half of those observed in our sample. A likely explanation for the minor difference may lie in grocery store margins excluding supplier mark-ups. Draganska, Klapper, and Villas-Boas (2010) find a rough 50/50 split in mark-ups between retailers and manufacturers in the German market for ground coffee, using scanner data from a national sample of stores belonging to six major retail chains. Portuguese and U.S. grocery store margins may also differ. In this sense, we are well within the expected interval for double marginalisation.

A significant point in favour of a lower-mark-up estimate is that the observed U.S. grocery store margins also follow the price index in Figure 2 well, more so when said index is adjusted to a four-quarter moving average. The revenue-weighted mean price index sees the

[12] Data from U.S. Census Bureau's Quarterly Financial Report (QFR) for Retail Trade, reproduced from a report by the White House Council of Economic Advisors, obtained at: https://bidenwhitehouse.archives.gov/cea/written-materials/2024/06/20/update-grocery-price-inflation-has-cooled-substantially.

[13] The data which does exist is annual, see Banco de Portugal's Quadros do Setor in https://www.bportugal.pt/QS/qsweb/Dashboards. It roughly corroborates the U.S. trends.

same doubling of prices between late 2020/early 2021 and late 2022/early 2023. This degree of implied cost pass-through matches a low-mark-up environment as estimated in our sample.

Beyond the difference in pricing assumption, other arguments can help reconcile the rest of the literature with the reported margins in our sample and at U.S. grocery stores. Single-category research has in the past targeted categories where supplier market power is especially evident, either due to market share concentration (e.g. soda, mark-ups estimated at between 19.9 and 63.9%; Dubé, 2005) or brand proliferation by large multi-product firms (e.g. ready-to-eat cereals, mark-up estimated at 35.8% under assumption of single-product firms; Nevo, 2001). The sample in Dopper et al. (2024) ends in 2019 and observes stable prices, while our sample starts in early 2020 and contains significant price volatility.

To obtain further insight on the distribution and significance of the mark-up trends, Figure 5 provides robust confidence intervals on the revenue-weighted mean mark-up estimates across mark-up percentiles.

Figure 5: Revenue-weighted mean mark-ups - percentiles

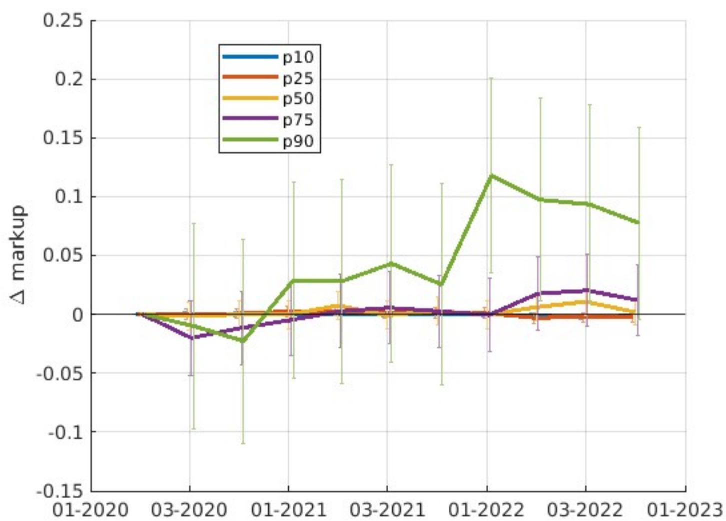

Mark-up changes remain insignificant for most goods across the sample period, except for a 10% increase in the first quarter of 2022 amongst the highest percentile group, which carries on for two more quarters before dissipating. This corresponds to the exact period where nominal prices separate from the CPI, likely due to supply chain disruption and an international energy shock. Notably, the increase in mark-ups observed in 2022 is not driven by supermarket heterogeneity, as can be seen in Figure 6.

Figure 6: Revenue-weighted mean mark-ups - supermarkets

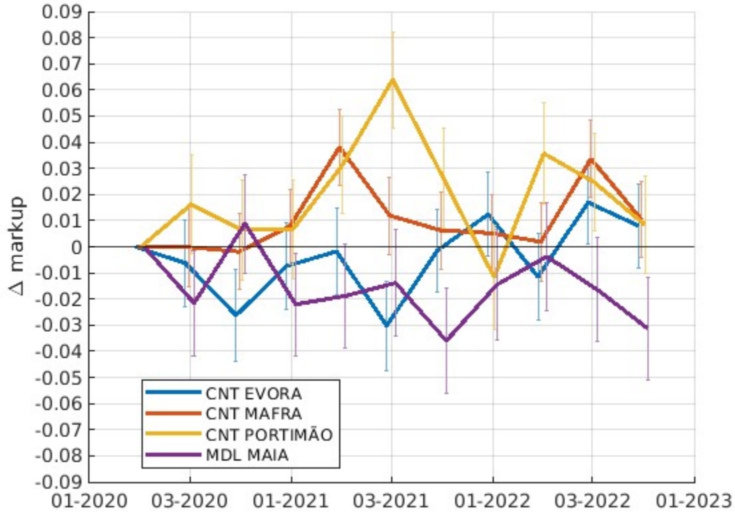

Significant heterogeneity is observed elsewhere, particularly in mid-2021, possibly due to a staggered propagation of the second COVID wave through Portugal. However, the majority of supermarkets face a similar increase in mark-ups in the third quarter of 2022.

To explore the causes of this late jump in market power, I consider a few measures of product mobility across the mark-up distribution in Figure 7. First, I recompute percentile bins within each category every quarter-year period, assign each product to a bin, and track transitions across adjacent quarter-years. A transition heatmap maps cross-time transitions, while a companion line plot shows the time path of said transitions.

Figure 7: Mark-up transition dynamics

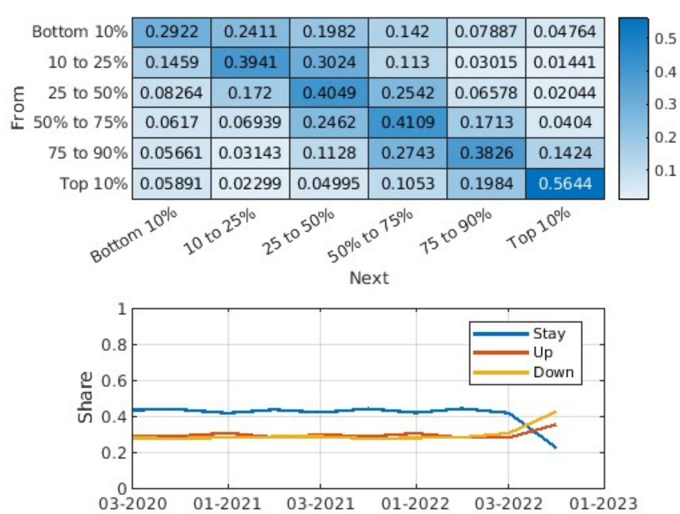

Overall, we observe substantial cross-time persistence, coupled with some degree of mean reversion. Products at the top end of the mark-up distribution, which we have noticed above to be a key driver of the overall mark-up increase observed at the end of our sample period, are especially persistent in their positioning. However, as can be seen in the line plot, persistence is shaken around the time where we have elsewhere observed a rise in market power, in mid-to-late 2022. This suggests substantial changes in tastes, with new goods rising to the top of the mark-up distribution as others fell.

To explore this further, in Figure 8 we observe Laspeyres- and Paasche-type fixed-basket mark-up indices, based on initial and final revenue weights respectively.

Figure 8: Fixed-Basket Mark-up Indices

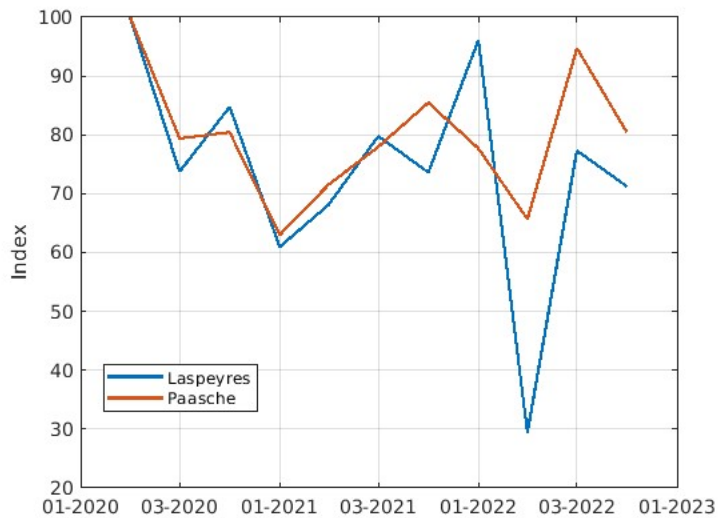

Divergence between them diagnoses reallocation effects; and in fact, as of the third quarter of 2022, the Paasche index sat above the Laspeyres index, revealing shifting preferences. This reflects reports from industry players, who suggest two key factors as playing an important role in this change in behaviour: (i) brand switching as a result of stockouts; and (ii) a move away from in-person dining, which as of 2024 had yet to rebound.[14]

8 Discussion

Using the set of observed unique transactions to represent A, rather than the true, latent consideration set, means that the above yields a still imperfect estimate of price elasticities. It is very likely we are overestimating the number of standard basis vectors missing; far more goods are likely considered standalone than we may observe even in a large sample. It is also

[14]See e.g. https://www.mckinsey.com/industries/consumer-packaged-goods/our-insights/state-of-consumer, and https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/survey-us-consumer-sentiment-during-the-coronavirus-crisis.

likely that the number of columns we have included in A is unnecessarily excessive to express its true span, raising computational concerns in some setups. Our approach strictly dominates doing nothing about corner solutions and is both testable and expandable via inclusion of more transactions data. However, our discussion of the implications of the structure of A for competitive outcomes may help us inform adjustments. I make two suggestions for others considering the implementation of the approach used here.

First, find all goods that ever appear alone - i.e. the set of singletons. For every unique transaction, remove any singleton goods from it and keep only the remaining goods; if nothing is left, drop that transaction because it can be rebuilt by adding singletons. Next, group transactions by the remaining goods they have in common. Within each group, if there is a transaction that contains exactly those remaining goods and no singletons, keep exactly that transaction and drop the others in the group. If there is no such transaction, keep only the transactions whose added singletons are minimal, meaning you drop any transaction whose added singletons strictly contain those of another kept basket; if two unique transaction's added singletons are different and neither contains the other, keep both. This keeps the same non-negative combinations that define the NN-LF as before, so our results don't change.

Second, we may add a simple, conservative screen to flag goods that are never bought on their own to a point that it is statistically unlikely we will observe such purchases in any larger sample. For each good, we look at how many times it appears in the data and how many of those times it appears by itself. Using a neutral Bayesian prior (e.g. Jeffreys) as the singleton rate we form a one-sided, credible upper-bound on that rate at a chosen confidence level. If a good is never seen alone and this bound is below a very small threshold (say a tenth of a percent), we may label it as never considered as a singleton basket with high confidence. A handy rule of thumb is the "rule of three": if you never see a singleton in N purchases, the 95% upper bound is roughly three over N, so large N gives strong evidence the true singleton rate is tiny. We can consider a few confidence levels and thresholds to see how many goods meet increasingly strict standards, then run diagnostics and robustness checks. This step does not change the NNLS geometry or the core estimation; it simply highlights goods that are, by the data, mostly likely considered standalone even if they are not observed as such.

9 Conclusion

In ignoring the relevance of co-purchase dependencies, textbook consumer choice understates market power and misclassifies rivals and complements. Simplified practical implementations, unconstrained and missing cross-category price effects, also misclassify market structure, but overestimate the actual markups.

As a whole, pursuing a realistic treatment of consumers' purchasing behaviour matters for inference and policy: when consideration binds, own-price effects attenuate, and cross-effects reconfigure. By raising complementarities and weakening own-price slopes when goods are co-considered, settings with choice bounded by consideration sets can imply greater consumer welfare under market concentration. This has major implications for merger policy.

Using a scalable linear-demand estimator that marries a co-purchase network proxy with a series expansion and a single Hausman-type competitor-price instrument regularized via PCA, I show that revenue-weighted mark-ups in select Portuguese grocery stores averaged about 11%, trough at $\approx 8.5\%$ in early-2021 and peak at $\approx 15\%$ in Q3-2022, closely tracking pandemic timing and later energy/logistics shocks with strong pass-through. The results mirror survey data on grocery store margins and suggest shifting tastes might have driven a small rise in market power at the end of the pandemic. The reliance of my approach on accessible, free-to-use point-of-sale data directly from a supermarket chain, rather than on granular scanner data from an external provider, makes it a valuable tool for researchers and practitioners wanting to explore industries for which such providers are either unavailable, or inaccessible.

Lastly, this paper's framework yields operational payoffs useful for day-to-day retail decisions, namely closed-form stock-out counterfactuals for prices, quantities, profits, and surplus.

Further research on endogenising the consideration set is desirable; it is likely that, beyond minor price movements and a stable product assortment in the short run, consideration set exogeneity and stability may be too restrictive a set of assumptions. Additional research is also needed to better understand the implications of large market asymptotics [15]; and test the predicted attenuation and complementarity patterns around natural shocks (e.g. VAT changes, packaging/size reforms, exit and entry).

10 References

Allen, R., and Rehbeck, J. (2022). Latent complementarity in bundles models. Journal of Econometrics, 228(2), 322-341.

Amir, R., Erickson, P., and Jin J. (2017). On the microeconomic foundations of linear demand for differentiated products. Journal of Economic Theory, 169, 641-665.

Armstrong, T. (2016). Large market asymptotics for differentiated product demand estimators with economic models of supply. Econometrica, 84(5), 1961-1980.

Armstrong, M., and Vickers J. (2015). Which demand systems can be generated by discrete choice? Journal of Economic Theory, 158(A), 293-307.

Atalay, E., Frost, E., Sorensen, A., Sullivan, C., and Zhu, W. (2023). Scalable demand and markups (NBER Working Paper). National Bureau of Economic Research.

[15] Otsu and Sunada (2024) suggest that demand shifters used as price instruments become weak where total market size is fixed and the number of products grows. This is a restrictive, specific asymptotic path that is not clearly justifiable or testable.

Berry, S., Khwaja, A., Kumar, V., Musalem, A., Wilbur, K., Allenby, G., Anand, B., Chintagunta, P., Hanemann, W., Jeziorski, P., and Mele, A. (2014). Structural models of complementary choices. Marketing Letters, 25(3), 245-256.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. Econometrica, 63(4), 841–890.

Blundell, R., and Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. Journal of Econometrics, 87(1), 115–143.

Bos, I., and Vermeulen, D. (2022). On the Microfoundation of Linear Oligopoly Demand. The B.E. Journal of Theoretical Economics, 22(1), 1-15.

Caplin, A., Dean, M., and Martin, D. (2011). Search and Satisficing. American Economic Review, 101(7), 2011.

Chen Y., and Riordan, M. (2007). Price and variety in the spokes model. The Economic Journal, 117(522), 897-921.

Choné, P., and Linnemer, L. (2020). Linear demand systems for differentiated goods: Overview and user's guide. International Journal of Industrial Organization, 73, 1–18.

Christensen, L., Jorgenson, D., and Lau, L. (1975). Transcendental Logarithmic Utility Functions. American Economic Review, 65(3), 367-383.

Deaton, A., and Muellbauer, J. (1980). An almost ideal demand system. American Economic Review, 70, 312-326.

Dixit, A., and Stiglitz, J. (1977). Monopolistic Competition and Optimum Product Diversity. American Economic Review, 67(3), 297-308.

Dopper, H., MacKay, A., Miller, N., and Stiebale, J. (2024). Rising Markups and the Role of Consumer Preferences. Journal of Political Economy, 133(8).

Draganska, M., Klapper, D., and Villas-Boas, S. (2010). A Larger Slice or a Larger Pie? An Empirical Investigation of Bargaining Power in the Distribution Channel. Marketing Science, 9(1), 57-74.

Dubé, J. (2005). Product Differentiation and Mergers in the Carbonated Soft Drink Industry. Journal of Economics & Management Strategy, 14(4), 879-904.

Dubois, P., Griffith, R., and Nevo, A. (2014). Do prices and attributes explain international differences in food purchases? American Economic Review, 104(3), 832–67.

Dubois, P., Griffith, R., and O'Connell, M. (2022). The use of scanner data for economics research. Annual Review of Economics, 14, 723-745.

Eliaz, K., and Spiegler, R. (2011). Consideration Sets and Competitive Marketing. Review of Economic Studies, 78(1), 235–262.

Ershov, D., Laliberté, J., Marcoux, M., and Orr, S. (2024). Estimating Complementarity with Large Choice Sets: an Application to Mergers. Available at SSRN: https://ssrn.com/abstract=3802097

Fan, Y. (2013). Ownership consolidation and product characteristics: A study of the US daily newspaper market. American Economic Review, 103(5), 1598–1628.

Fosgerau, M., Monardo, J., and de Palma, A. (2024). The Inverse Product Differentiation Logit Model. American Economic Journal: Microeconomics, 16(4), 329-370.

Fox, T., and Lazzati, N. (2017). A note on identification of discrete choice models for bundles and binary games. Quantitative Economics, 8(3), 1021-1036.

Gandhi, A. and Houde, J. (2019). Measuring substitution patterns in differentiated-products industries. NBER Working Paper, 26375.

Gandhi, A., and Nevo, A. (2021). Empirical models of demand and supply in differentiated products industries. In K. Ho, A. Hortaçsu, and A. Lizzeri (Eds.), Handbook of Industrial Organization, 4(1), 63–139. Elsevier.

Gentzkow, M. (2007). Valuing new goods in a model with complementarity: Online newspapers. American Economic Review, 97(4), 713-744.

Gorman, W. (1959). Separable Utility and Aggregation. Econometrica, 27(3), 469-481.

Hauser, J. (2014). Consideration-set Heuristics. Journal of Business Research, 67(8), 1688-1699.

Hauser, J. and Wernerfelt, B. (1990). An Evaluation Cost Model of Consideration Sets. The Journal of Consumer Research, 16(4), 393-408.

Hausman, J. (1996). Valuation of new goods under perfect and imperfect competition. In: The Economics of New Goods, 207-248. National Bureau of Economic Research, Inc.

Hausman, J., Leonard, G., and Zona J. (1994). Competitive analysis with differentiated products. Annales d'Économie et de Statistique, 34, 159-180.

Hurewicz, W., and Wallman, H. (1941). Dimension theory. Princeton University Press.

Iaria, A., and Wang, A. (2019). Identification and Estimation of Demand for Bundles. Available at SSRN: https://ssrn.com/abstract=3458543

Lanier, J., Large, J., and Quah, J. (2023). Estimating very large demand systems. Available at https://ora.ox.ac.uk/objects/uuid:13accd61-5e96-46e8-849a-fdd7c627cbfd

Lawson, C., and Hanson, R. (1974). Solving Least Squares Problems. Prentice-Hall, Englewood Cliffs, N.J.

Lee, Y. (2024). Shrinkflation: Evidence on Product Downsizing and Consumer Response. Kilts Center at Chicago Booth Marketing Data Center Paper, Available at SSRN: https://ssrn.com/abstract=5053745 or http://dx.doi.org/10.2139/ssrn.5053745.

Lee, L., and Pitt, M. (1986). Microeconometric Demand System with Binding Nonnegativity Constraints: The Dual Approach. Econometrica, 54(5), 1237-1242.

Lewbel, A., and Nesheim, L. (2019). Sparse demand systems: corners and complements. Boston College Working Papers in Economics 1005, Boston College Department of Economics.

Magnolfi, L., McClure, J., and Sorensen, A. (2025). Triplet embeddings for demand estimation. American Economic Journal: Microeconomics, 17(1), 282–307.

Manchanda, P., Ansari, A., and Gupta, S. (1999). The "Shopping Basket": A Model for Multicategory Purchase Incidence Decisions. Marketing Science, 18(2), 95-114.

Manzini, P., Mariotti, M., and Ülkü, L. (2019). Stochastic Complementarity. The Economic Journal, 129(619), 1343-1363.

Matějka, F., and MacKay, A. (2015). Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. American Economic Review, 105(1), 272–98.

McFadden, D. (2001). Economic Choices. American Economic Review, 91(3), 351-378.

McFadden, D. (1973) Conditional Logit Analysis of Qualitative Choice Behavior. In: Zarembka, P., Ed., Frontiers in Econometrics, Academic Press, 105-142.

Moore, E. (1920). On the reciprocal of the general algebraic matrix. Bulletin of the American Mathematical Society, 26(9), 394–395.

Nevo, A. (2011). Empirical Models of Consumer Behavior. Annual Reviews, 3, 51-75.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. Econometrica, 69(2), 307–342.

Nocke, V., and Schutz, N. (2025). Oligopoly, Complementarities, and Transformed Potentials. CEPR Discussion Paper No. 20527. CEPR Press, Paris & London.

Otsu, T., and Sunada, K. (2024). On large market asymptotics for spatial price competition models. Economics Letters, 234.

Pellegrino, B. (2025). Product differentiation and oligopoly: A network approach. American Economic Review, 115(4), 1170–1225.

Penrose, R. (1955). A generalized inverse for matrices. Proceedings of the Cambridge Philosophical Society. 51(3), 406–413.

Pinkse, J., and Slade, M. (2004). Mergers, brand competition, and the price of a pint. European Economic Review, 48(3), 617-643.

PORDATA. (2025). Contemporary Portugal Database. Retrieved from http://www.pordata.pt.

Ruiz, F., Athey, S., and Blei, D. (2020). SHOPPER: A probabilistic model of consumer choice with substitutes and complements. The Annals of Applied Statistics, 14(1), 1-27.

Simon, H. (1957). Models of Man: Social and Rational. New York: John Wiley and Sons.

Spence, M. (1976). Product Selection, Fixed Costs, and Monopolistic Competition. Review of Economic Studies, 43(2), 217-235.

Theil, H. (1965). The Information Approach to Demand Analysis. Econometrica, 33(1), 67-87.

Tian, Y., Lautz, S., Wallis, A., and Lambiotte, R. (2021). Extracting complements and substitutes from sales data: A network perspective. EPJ Data Science, 10(45).

Tirole, J. (1988). The Theory of Industrial Organization. Cambridge, MA: The MIT Press.

Vives, X. (1987). Small Income Effects: A Marshallian Theory of Consumer Surplus and Downward Sloping Demand. Review of Economic Studies, 54(1), 87-103.

Wales, T., and Woodland, A. (1983). Estimation of consumer demand systems with binding non-negativity constraints. Journal of Econometrics, 21(3), 263-285.

Wright P., and Barbour, F. (1977). Phased Decision Strategies: Sequels to an Initial Screening. Stanford School of Business Working Paper No. 353.

A Appendix

A.1 Proofs

Lemma 1. By construction, $\boldsymbol{x}'M\boldsymbol{x} > 0$ for all non-zero $\boldsymbol{x}$. $\boldsymbol{y}'A'MA\boldsymbol{y} \geq 0$ because, where $A\boldsymbol{y} \neq 0$, $\boldsymbol{y}'A'MA\boldsymbol{y} = \boldsymbol{x}'M\boldsymbol{x} > 0$ for non-zero $\boldsymbol{x}$. If $A\boldsymbol{y} = 0$, $\boldsymbol{x}'M\boldsymbol{x} = 0$ because $\boldsymbol{x} = A\boldsymbol{y} = 0$. $\square$

Lemma 2. Regarding symmetry: for M positive definite and symmetric (by construction), $A'MA$ is symmetric - $(A'MA)' = A'M'A = A'MA$. By eigenvalue decomposition, $A'MA = U\Lambda U'$, and $(A'MA)^+ = U\Lambda^+U'$, where Λ^+ is formed by replacing each non-zero $\lambda, \forall i$ with $1/\lambda_i$, and leaving zeros intact. Since Λ^+ is diagonal, it is symmetric, and therefore $(A'MA)^+$ is also symmetric. Lastly, if $(A'MA)^+$ is symmetric, $A(A'MA)^+A'$ is also symmetric - $(A'(A'MA)^+A')' = A'((A'MA)^+)'A' = A(A'MA)^+A'$. Regarding positive semi-definiteness: If $A'MA$ is positive semi-definite, $(A'MA)^+$ is also positive semi-definite. As shown above, $A'MA$ is symmetric and can be diagonalised via eigenvalue decomposition, and the pseudoinverse inverts the non-zero eigenvalues while leaving the zero eigenvalues at zero. The inversion does not change the signs of the eigenvalues, so $\forall\lambda \geq 0$ still holds. If $(A'MA)^+$ is positive semi-definite, $\boldsymbol{x}'(A'MA)^+\boldsymbol{x} \geq 0$. This is the case for $\boldsymbol{x} = A'\boldsymbol{y} \neq 0$. For $\boldsymbol{x} = A'\boldsymbol{y} = 0$, $\boldsymbol{x}'(A'MA)^+\boldsymbol{x} = 0$. Therefore, $\boldsymbol{x}'(A'MA)^+\boldsymbol{x}' = \boldsymbol{y}'A(A'MA)^+A'\boldsymbol{y} \geq 0$ for all non-zero (and zero) $\boldsymbol{y}$. $\square$

Lemma 3. The sum of a positive definite with a positive semi-definite matrix is not always invertible. By theorem, however, the sum of a positive definite Ω with a positive semi-definite matrix $A'(A'MA)^+A'$, *where the positive definite matrix is diagonal with strictly positive diagonal elements*, is itself positive definite, and therefore invertible: $\boldsymbol{x}'\Omega + A(A'MA)^+A'\boldsymbol{x} = \boldsymbol{x}'\Omega\boldsymbol{x} + \boldsymbol{x}'A(A'MA)^+A'\boldsymbol{x} = \text{positive} + \text{non-negative} > 0$. We can show that Ω is strictly positive in its diagonal. Express the diagonal matrix whose diagonal matches that of $A(A'MA)^+A'$ as Ω. To guarantee strictly positive elements in the diagonal of Ω, two requirements must be satisfied. Firstly, $\boldsymbol{v}_i \neq 0, \forall i$, for $\boldsymbol{v}_i = A'\boldsymbol{e}_i$, where $\boldsymbol{e}_i$ is the standard basis vector, and $\Omega_{ii} = \boldsymbol{e}_i'A(A'MA)^+A'\boldsymbol{e}_i = \boldsymbol{v}_i'(A'MA)^+\boldsymbol{v}_i$, where $\boldsymbol{e}_i$ "picks out" the i-th diagonal of Ω. This is the case as long as row i of A has at least one non-zero term, which is true of every row in A by definition, as otherwise said row would describe a product we have not observed in the data. Secondly, $\boldsymbol{v}_i'(A'MA)^+\boldsymbol{v}_i > 0, \forall i$. Now, as we mentioned earlier, $(A'MA)^+$ is positive semi-definite. However, $\boldsymbol{v}_i$ is restricted to the column space of A by construction, the same column space of $(A'MA)^+$. On the subspace where $A'MA$ is "active" (its column space), the pseudoinverse $(A'MA)^+$ behaves like a true inverse and is positive definite on that subspace. This ensures that whenever is $\boldsymbol{v}_i$ is nonzero (i.e., when it lies in the column space of A), we have $\boldsymbol{v}_i'(A'MA)^+\boldsymbol{v}_i > 0, \forall i$. Given our proof of strictly positive diagonal elements in Ω, $\Omega + A(A'MA)^+A'$ is invertible. $\square$

Lemma 4. All model outputs are in one of the following forms: $\boldsymbol{y_1} = \alpha\boldsymbol{x}'\Sigma\boldsymbol{x}$ or $\boldsymbol{y_2} = \alpha\Sigma\boldsymbol{x}$ for α a strictly positive scalar, $\boldsymbol{x}$ a strictly positive vector, and Σ a positive semi-definite matrix. Suppose we wish to consider the impacts derived from the removal of the i-th component of $\boldsymbol{x}$ and Σ. These terms can be expressed as such:

$$\boldsymbol{x} = \begin{pmatrix} x_i \\ \boldsymbol{x}_{-i} \end{pmatrix} \tag{45}$$

and

$$\Sigma = \begin{pmatrix} \sigma_{ii} & \boldsymbol{\sigma}_i' \\ \boldsymbol{\sigma}_i & \Sigma_{-i} \end{pmatrix} \tag{46}$$

The notation for $\boldsymbol{\sigma}_i$ excludes the element corresponding to σ_{ii}. Starting with $\boldsymbol{y_1}$, we may restate it as:

$$\boldsymbol{y_1} = \sigma_{ii}x_i^2 + 2x_i(\boldsymbol{\sigma}_i'\boldsymbol{x}_{-i}) + \boldsymbol{x}_{-i}\Sigma{-i}\boldsymbol{x}_{-i} \tag{47}$$

of which the first two terms are the contribution from the i-th dimension of both $\boldsymbol{x}$ and Σ. If we remove the i-th product, i.e. the i-th dimension of $\boldsymbol{x}$ and Σ, the difference between the original and reduced term is:

$$\Delta_1 = -\alpha\big[\sigma_{ii}x_i^2 + 2x_i(\boldsymbol{\sigma}_i'\boldsymbol{x}_{-i})\big] \tag{48}$$

We can similarly re-write $\boldsymbol{y_2}$:

$$\boldsymbol{y_2} = \alpha\begin{pmatrix} \sigma_{ii}x_i + \boldsymbol{\sigma}_i'\boldsymbol{x}_{-i} \\ \boldsymbol{\sigma}_i x_i + \Sigma_{-i}\boldsymbol{x}_{-i} \end{pmatrix} \tag{49}$$

If the i-th component is dropped, the first (upper) block drops out, and the lower block's reduced version is $\alpha\Sigma_{-i}\boldsymbol{x}_{-i}$. □

Lemma 5. If A is full row rank, it can be shown that $A(A'MA)^{+}A' = M^{-1}$. Its singular value decomposition could be written as $A = U\Sigma V'$, for U a $K \times K$ orthogonal matrix; $\Sigma = [\Sigma_r 0]$ a $K \times J$ matrix, with Σ_r a $K \times K$ diagonal matrix containing the nonzero singular values of A, and 0 a $K \times (J-K)$ zero block; and $V = [V_1 V_2]$ is a $J \times J$ orthogonal matrix where V_1 is a $J \times K$ matrix whose columns form the orthonormal basis for the row space of A, while V_2 is a $J \times (J-K)$ matrix spanning the null space of A. A can then be re-written as $A = U\Sigma_r V_1'$. From here, $A(A'MA)^{+}A' = U\Sigma_r V_1'(V_1\Sigma_r U'MU\Sigma_r V_1')^{+}V_1\Sigma_r U' = U\Sigma_r V_1'V_1(\Sigma_r U'MU\Sigma_r)^{-1}V_1'V_1\Sigma_r U' = U\Sigma_r(\Sigma_r U'MU\Sigma_r)^{-1}\Sigma_r U' = U(U'MU)^{-1}U' = M^{-1}$ by the properties of orthogonal and diagonal matrices. □

Lemma 6. For a given M, define the M-norm: $||x||_M = \sqrt{x'Mx}$. For a given price change $\Delta\boldsymbol{p}$, $\Delta\boldsymbol{q}_{miss} = \phi M^{-1}\Delta\boldsymbol{p}$ and $\Delta\boldsymbol{q}_{corr} = \phi A(A'MA)^{+}A'\Delta\boldsymbol{p}$ for the misspecified and correct model respectively. It can be shown that $||[M^{-1} - A(A'MA)^{+}A']\Delta\boldsymbol{p})||_M \geq 0$. To see how, note that for $W = M^{1/2}A$, $A(A'MA)^{+}A' = M^{-1/2}PM^{-1/2}$, for $P = W(W'W)^{+}W'$. P is the orthogonal projector onto the column-space of W; it is idempotent and symmetric. We can re-write $M^{-1} - A(A'MA)^{+}A' = M^{-1/2}(I-P)M^{-1/2}$, where $I-P$ is the orthogonal projector onto the complement of $col(W)$ (and also idempotent and symmetric). Orthogonal projections are positive semi-definite. Lastly, it is trivial to show that $||[M^{-1} - A(A'MA)^{+}A']\Delta\boldsymbol{p})||_M^2 \geq 0$, from where $||[M^{-1} - A(A'MA)^{+}A']\Delta\boldsymbol{p})||_M \geq 0$ follows naturally. Now, it is a well known result that all norms on $\mathbb{R}^{m\times n}$ are *equivalent* (Golub and Van Loan, 2013, pp. 73), which is to say that there exist constants $c, C > 0$ such that, for every x in a given subspace, $c||x||_a < ||x||_b < C||x||a$. Therefore, $||[M^{-1} - A(A'MA)^{+}A']\Delta\boldsymbol{p}||_M \geq 0 \Leftrightarrow ||[M^{-1} - A(A'MA)^{+}A']\Delta\boldsymbol{p}||_1 \geq 0$. Note that the inequality holds strictly if $\Delta\boldsymbol{p} \notin col(A)$, and is equal to zero otherwise. For a given subspace (e.g. that defined by $col(A)$), every $\Delta\boldsymbol{p}$ can be decomposed as follows (Golub and Van Loan, 2013, pp. 82): $\Delta\boldsymbol{p} = P\Delta\boldsymbol{p} + (I-P)\Delta\boldsymbol{p} = \Delta\boldsymbol{p}_{col(A)} + \Delta\boldsymbol{p}_{null(A)}$ for P the (unique) orthogonal projection onto the subspace. Every price shock $\Delta\boldsymbol{p} \notin col(A)$ loses $\Delta\boldsymbol{p}_{null(A)}$ and is affected, even if it pertains to products not directly involved in the linear dependency.[16] $\square$

Lemma 7. The intermediate step is to show that $A(A'MA)^{+}A'M = M^{-1}A(A'M^{-1}A)^{+}A'$; if this is the case, both sides of the equation can be multiplied by M^{-1} to yield the Lemma's equality. I do so by noting that both sides of the equation are an M-orthogonal projection onto the same space $col(A)$. Since orthogonal projections are unique, the two must be equivalent. An orthogonal projection to space $col(A)$ must satisfy three conditions: the matrix must be in $col(A)$, be idempotent, and self-adjoint in M.

For a vector $\boldsymbol{u} = A\boldsymbol{v} \in col(A)$, $A(A'MA)^{+}A'M(Av) = Av$. Similarly, noting that $col(M^{-1}A) = col(A)$ as M^{-1} is positive definite with $col(M^{-1}) = \mathbb{R}$, we have $\boldsymbol{u} = M^{-1}A\boldsymbol{v} \in col(A)$ such that $M^{-1}A(A'M^{-1}A)^{+}A'(M^{-1}Av) = M^{-1}Av$. This means that every vector in $col(A)$ is also in the column space of these expressions. The proof is complete by noting the reverse - the column space of these expressions is also a subset of $col(A)$: $A(A'MA)^{+}A'M = AB$ for $B = (A'MA)^{+}A'M$, and $col(AB) \subseteq col(A)$; $M^{-1}A(A'M^{-1}A)^{+}A' = M^{-1}AB$ for $B = (A'M^{-1}A)^{+}A'$, with $col(M^{-1}AB) \subseteq col(M^{-1}A) = col(A)$. Idempotency is similarly straightforward: $A(A'MA)^{+}A'M \cdot A(A'MA)^{+}A'M = A(A'MA)^{+}A'M$ and $M^{-1}A(A'M^{-1}A)^{+}A' \cdot M^{-1}A(A'M^{-1}A)^{+}A' = M^{-1}A(A'M^{-1}A)^{+}A'$. Lastly, note:

$$(A(A'MA)^{+}A'M\boldsymbol{x})'M\boldsymbol{y} = \boldsymbol{x}'M(A(A'MA)^{+}A'M\boldsymbol{y}) \tag{50}$$

and

$$(M^{-1}A(A'M^{-1}A)^{+}A'\boldsymbol{x})'M\boldsymbol{y} = \boldsymbol{x}'M(M^{-1}A(A'M^{-1}A)^{+}A'\boldsymbol{y}) \tag{51}$$

$\square$

[16] How likely are $\Delta\boldsymbol{p} \in col(A)$? The answer is *very unlikely.* A well known result in measure theory states that any linear sub-space of less than full rank - in this case rank less than K - has a K-dimensional Lebesgue measure 0. By implication, any K-dimensional price shock $\Delta\boldsymbol{p}$ drawn from a continuous distribution - i.e. no draw has positive probability - will have a probability zero of lying in $col(A)$.

A.2 Robustness of the Representative Consumer Assumption

Quadratic utility models offer a compelling combination of analytical tractability and interpretation, yet they have long been criticised for their aggregation properties when consumers exhibit heterogeneous preferences. In the section that follows, we extend the representative consumer framework to incorporate such heterogeneity, addressing both the technical and conceptual concerns that arise when individual demand functions are aggregated. The key insight is that when each consumer's preferences are cast in quadratic form-specifically, when every consumer considers every available good-the model naturally aligns with Gorman polar form. This structure guarantees that, even with individual variations, the average demand preserves the same functional form as in the homogeneous case.

Assuming $\boldsymbol{q}_k > 0,\ \forall k \in \mathcal{K}$ for each consumer $n = 1, \ldots, N$, individual demand across goods is:

$$\boldsymbol{q}^*(\boldsymbol{p}, n) = A_n(A_n' M_n A_n)^+ A_n'(\boldsymbol{\delta} + \phi \boldsymbol{p}) \tag{52}$$

and total demand is

$$\boldsymbol{q}^*(\boldsymbol{p}) = \sum_{n=1}^{N} \boldsymbol{q}^*(\boldsymbol{p}, n) \tag{53}$$

Because QQUM is in Gorman polar form, the aggregation is exact and total demand can be restated as

$$\boldsymbol{q}^*(\boldsymbol{p}) = N \cdot A(A'MA)^+ A'(\boldsymbol{\delta} + \phi \boldsymbol{p}) \tag{54}$$

for

$$A(A'MA)^+ A' = \frac{1}{N} \sum_{n \in N} A_n(A_n' M_n A_n)^+ A_n' \tag{55}$$

We have assumed homogeneous $\boldsymbol{\delta} + \phi \boldsymbol{p}$ for simplicity. For heterogeneous $\boldsymbol{\delta(n)} + \phi(n)\boldsymbol{p} > 0,\ \forall n$, true by construction, this marginal utility term also averages out. The representative consumer model described throughout this paper thus approximates average demand across heterogeneous consumers, due to Gorman-aggregability.

Implicit in the assumption that $\boldsymbol{q}_k > 0, \forall k \in \mathcal{K}$ across all consumers is the idea that every consumer's consideration set includes every available good. But suppose instead that the goods present across consumers' consideration sets differ. For example, let $K - L$ be the number of goods that all consumers (the "common" goods) have in their consideration set, and let L be the number of "special" goods that only a subset of consumers consider. We can still maintain a QQUM specification for each consumer by collapsing the L special goods into a single composite good S, for which

$$q_S = \sum_{l=1}^{L} \omega_l q_l \tag{56}$$

with positive weights ω_l. Now, rather than keeping separate quantities for each of the K goods, we represent the full consumer bundle as $(\boldsymbol{q}_{-S}, q_S)$. Because QQUM is already in Gorman polar form, the heterogeneity that might arise (including differences in substitution patterns among the "special" goods) is entirely captured in the quadratic term. In other words, all complementarity and substitution patterns appear in the M matrix. When we collapse the special goods, we must aggregate the corresponding elements of M as follows. Suppose we partition the original M as:

$$M = \begin{pmatrix} M_{-S,-S} & M_{S,-S} \\ M_{-S,S} & M_{S,S} \end{pmatrix} \tag{57}$$

where $M_{-S,-S}$ pertains to the common goods, $M_{S,S}$ to the special goods, and $M_{S,-S} = M'_{-S,S}$ to the cross-terms. By collapsing the special goods into S, utility becomes:

$$U(\boldsymbol{q}) = \boldsymbol{q}'_{-S}\boldsymbol{\delta}_{-S} + \delta_S q_S - \frac{1}{2}\left(\boldsymbol{q}_{-S} \quad q_S\right) M^* \begin{pmatrix} \boldsymbol{q}_{-S} \\ q_S \end{pmatrix} + q_0 \tag{58}$$

with δ_{-S} the weighted average of the corresponding elements in $\boldsymbol{\delta}$ for the special goods, and M^* defined as:

$$M^* = \begin{pmatrix} M_{-S,-S} & M_{S,-S} \cdot \boldsymbol{\omega} \\ (M_{S,-S} \cdot \boldsymbol{\omega})' & \boldsymbol{\omega}' \cdot M_{S,S} \cdot \boldsymbol{\omega} \end{pmatrix} \tag{59}$$

i.e. each common good's cross-effect with the composite is a weighted sum of its effects on the special goods. In other words, by "collapsing" the L special goods into one composite good through a weighted sum, we are effectively aggregating their quadratic interactions. The implication for the M matrix is that its entries related to the special goods get averaged (via the weights ω) into a single scalar (for the composite good) and into cross terms with the common goods. If every consumer is modeled this way, and if the weights (or the heterogeneity in preferences over the special goods) average out appropriately, the aggregate demand function will retain the same structural form as in the fully common case.

This may however not be convincing for the more discrete-choice-inclined IO economist. I therefore reproduce here the Choné and Linnemer (2020) generalisation of the Armstrong and Vickers (2015) result on the consistency of linear demand with a discrete-choice model where valuations are support-restricted. Let ζ_{kk} be the mass of consumers interested uniquely in brand $k, k = 1, \ldots, K$, and ζ_{kh} be the mass of consumers potentially interested in both good k and $h \neq k$ with $\zeta_{kh} = \zeta_{hk}$. Demand q has the form:

$$q_{kk} = x_{kk} - \tau_{kk}p_k \text{ and } q_{kh} = \begin{cases} x_{kh} + \tau_{kh}(p_h - p_k) & \text{if } \upsilon_{kh} < 0 \\ x_{kh} + \tau_{kh}(p_h + p_k) & \text{if } \upsilon_{kh} > 0 \end{cases} \tag{60}$$

where $\tau_{kk} > 0$ (own-price effect must be negative) and $\tau_{kh} = \tau_{hk}$ is a function of the strength of the interaction effect between the two goods, positive where $\upsilon_{kh} < 0$ and that interaction is substitution, negative when $\upsilon_{kh} > 0$ and the goods are complements. In the former case, $0 < x_{kh} < 1$ is the demand for good k when $p_k = p_h$, with $x_{hk} = 1 - x_{kh}$; in the latter, $0 < x_{kh} = x_{hk}$ are the demand for both good k and h when $p_k = p_h = 0$. In this case, we can therefore express total demand for a given product k as:

$$q_k = \sum_h \zeta_{kh} q_{kh} = (\sum_h \zeta_{kh} x_{kh}) - (\sum_h \zeta_{kh}\tau_{kh})p_k + \sum_{h \neq k} \zeta_{kh}\tau_{kh}p_h \tag{61}$$

which can alternatively be defined as $q_k = d_k - \upsilon_{kk}p_k - \sum_{h\neq k} h_{kh}p_h$ for appropriate restrictions on Υ. What are these and how do they limit the model? By construction, we have implicitly set $\upsilon_{kh} = -\zeta_{kh}\tau_{kh}$ and $\upsilon_{kk} = \sum_{\upsilon} \zeta_{kh}\tau_{kh}$. Combining the two, we have that $\upsilon_{kk} = \zeta_{kk}\tau_{kk} - \sum_{h\neq k} \upsilon_{kh}$. Since it must be that $\zeta_{kk}\tau_{kk} > 0$, this results in a requirement that $\upsilon_{kk} + \sum_{h\neq k} \upsilon_{kh} > 0$, i.e. if the prices of all goods are increased by the same amount then the demands of all goods decrease. For as long as this condition, a weaker assumption than diagonal dominance but stronger than the Law of Demand, is satisfied, then the QQUM for a representative consumer can be micro-founded by a set of heterogeneous consumers distributed along a Chen and Riordan (2007)-like radial network representation of the consumer preference space with Hotelling-like utility functions for each pair of products.[17]

[17] For a further discussion of this approach, see Bos and Vermeulen (2022).

A.3 Participation rates

In our empirical application, all stores were at least 1km away from the nearest store competitor at the beginning of the sample in 2020, except for Portimão (700m away). Maia faced entry by a competitor 700m away from its store in February 2023, just after the end of our sample, and another 1.4km away in late February 2022; Mafra faced (re-)entry 300m away in February 2021; and Portimão faced an additional entry 600m away in June 2022. This high level of activity in the industry reflects a recent expansion strategy pursued by European discount retailers in the Portuguese market, which remained active throughout the COVID pandemic. Given this paper's focus on intra-supermarket competition, these may raise concerns that inter-supermarket competition impacts our outcomes.

In addition, in the standardised regressions above, all product quantities purchased are corrected by the number of transactions observed each period in the relevant supermarket. This is in line with linear demand aggregation, as can be seen in the earlier section on the robustness of the representative consumer. However, such correction introduces endogeneity concerns: if prices at the grocery store increase, consumers may decide to shop elsewhere, and vice versa. What may be the implications for mark-ups? Start from the differentiated-product, supermarket-time specific, Bertrand-Nash, Lerner index:

$$L_{m,t} = -((G \circ E_{m,t})^{-1}) \cdot \mathbf{1} \tag{62}$$

where G is the ownership matrix, and

$$E_{m,t} = J_{m,t} \cdot p_{m,t}/s_{m,t} \tag{63}$$

with s the per-customer demand, and $J = \frac{\partial s}{\partial p}$, each specified per market m and time period t. If participation N varies with prices, the per-customer Jacobian we estimated earlier from demand, $\hat{J}_{m,t} = \frac{\partial s}{\partial p}$, bundles two margins:

$$\frac{\partial s}{\partial p} = (1/N) \cdot \frac{\partial q}{\partial p} - s \cdot (\eta^N)' \tag{64}$$

where $\eta_j^N \equiv \frac{\partial \ln N}{\partial p_j}$ (participation sensitivity). Thus the intensive-margin Jacobian needed for mark-ups is:

$$J_{m,t}^{int} = \hat{J}_{m,t} + s_{m,t} \cdot (\eta_{m,t}^N)' \tag{65}$$

If $\frac{\partial q}{\partial p} < 0$ (higher prices reduce participation), ignoring participation would make own-price slopes look too small in magnitude, which would inflate the Lerner indices. Adding $s \cdot (\eta^N)'$ would be necessary to restore the intensive margin, typically lowering the Lerner index relative to the uncorrected approach.

In this section, I discuss the matter of participation rates and verify the robustness of our estimates of price elasticity of demand both to inter-supermarket competition and price endogeneity.

A.3.1 Transaction trends

Consider the time trends for the number of transactions observed per supermarket in the sample period:

Figure 9: Per-supermarket transaction trends

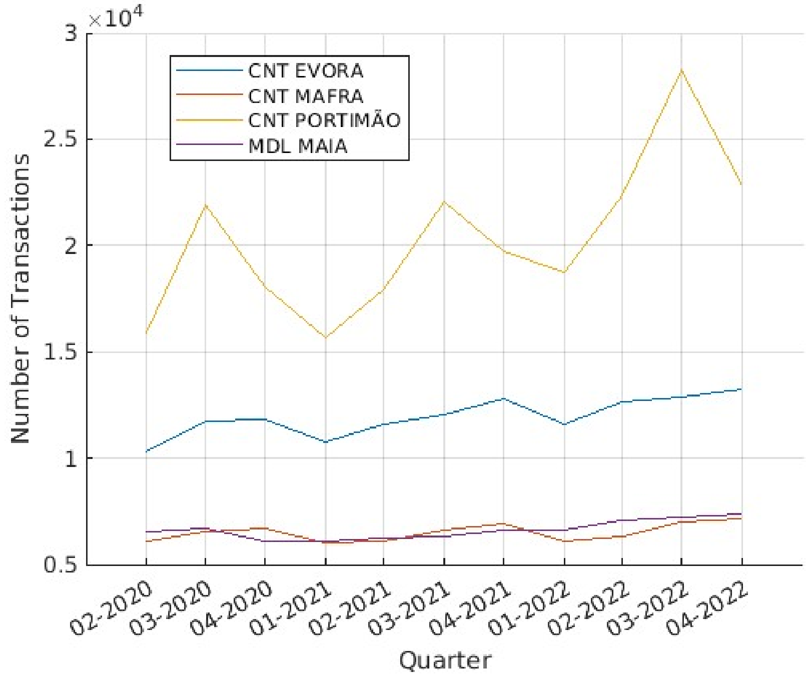

The trends do not resemble those observed for revenue-weighted mean prices in Figure 2. Nonetheless, the lack of a clear relationship between the data on grocery stores' prices and transactions numbers may mask a participation elasticity. Core inflation rose substantially in Portugal through the period of analysis. Nominal inflation pushed all prices up, but if as a result people shifted spending from elsewhere to groceries, this may have led to more transactions in-store. For example, the price of home food relative to restaurant meals may have fallen, attracting shoppers despite higher nominal grocery prices. This could be masking the correct price elasticities of grocery store purchases.

A.3.2 Estimation of participation rate elasticity

We wish to identify the causal effect of grocery store prices on participation N, net of common inflation and fixed market traits. The proposed estimation strategy is a first-difference IV panel estimator:

$$\Delta \ln(N_{m,t}) = \delta \cdot \Delta \ln P_{m,t} + \tau_t + \Delta u_{m,t} \tag{66}$$

Parameter δ is the participation elasticity. The number of transactions is calculated on a per-quarter, per-supermarket basis, while the price index $P_{m,t}$ is set under two specifications: (i) a simple average across goods; and (ii) a revenue-weighted function of per-quarter, per-supermarket goods prices. This approach matches our earlier empirical specification, which used quarterly data, but also limits the size of the regression we can run. Variable τ_t is a time trend.

Taking first differences removes latent time-invariant terms while preempting trend drifts. An instrument is necessary in this setting, as before, to isolate the price effect on participation rates with exogenous variation, avoiding reverse causality. I use the change in a good's competitors' average price (our previous instrument), aggregated into an index as that which it instruments for.

For robustness, I consider current, 1–2 lags and select that with the highest first-stage F. With one endogenous regressor, we run 2SLS and report Eicker–Huber–White standard errors; the 2SLS p-value; the Anderson–Rubin p-value, which stays valid even when instruments are weak; and the (homoskedastic) first-stage F for the chosen instrument set. Results are shown below:

Table 8: Participation Shares - IV Results by Specification

Price index	$\hat{\delta}$ (SE)	p-value	AR p-value	First-stage F	N
Simple average	0.056 (2.685)	0.983	0.981	6.920	32
Revenue-weighted	−2.855 (4.890)	0.559	0.223	6.119	32

Notes: $\hat{\delta}$ is the IV estimate; standard errors in parentheses. "AR p-value" is the Anderson–Rubin test p-value. "First-stage F" is the first-stage F-statistic. N is the number of observations.

The preferred instrument is the first lag of the competitor price change with a first-stage $F \approx 6.12$ and 6.92 respectively, still however indicating weak instrument strength. Nonetheless, neither specification shows a detectable effect of prices on participation, even under the Anderson–Rubin test. The simple-average specification yields $\hat{\delta} \approx 0.06$ and the revenue-weighted specification $\hat{\delta} \approx -2.86$. Deflating the price indices by the CPI leaves the results unchanged.

Though underpowered, the results suggest participation sensitivity is not a significant driver of our results. This seems to confirm that (i) inter-supermarket competition is not a major problem in the set of stores and period I have sampled, and (ii) adjusting demand by the number of transactions observed each period in each store raises no endogeneity concerns within our sample. Nonetheless, the approach used here could be averaging out what are likely important effects amongst staple goods in any supermarket assortment. The reader is therefore encouraged to perceive this paper's mark-up estimates as upper bounds, especially so for goods on the upper tail of the revenue distribution.