Approximating demand dynamics in antitrust policy

Alan Crawford[†]

March, 2019

Abstract

Empirical analysis for antitrust policy commonly uses accounting margins as inputs into merger screening tools. This paper shows how these margins can be combined with a static demand model to estimate a set of price elasticities for storable goods that are consistent with dynamic demand responses to permanent, rather than transitory price changes. As a result, demand dynamics that create inter-temporal substitution for storable goods are better captured by the resulting set of price elasticities. To illustrate this method, I apply it to the UK laundry detergent industry from 2002 to 2012. I present evidence that product innovations in this industry that lower storage costs affect demand dynamics and hence the degree of bias of price elasticities that come from static demand models. I also show how adjusting price elasticities to reflect demand dynamics can lead to different policy conclusions. I illustrate this by assessing whether there is any evidence that anti-competitive conduct of a laundry detergent cartel in mainland Europe had any effect on the UK laundry detergent market.

Acknowledgements: I would like to thank Lars Nesheim, Michela Tincani, Pedro Mira, Mateusz Mysliwski and Marleen Marra for helpful comments and suggestions. I would also like to thank the Economic Social Research Council (ESRC) for funding this research. I also gratefully acknowledge financial support from the European Research Council (ERC) under ERC-2009-AdG grant agreement 249529. Data supplied by Kantar Worldpanel. The use of Kantar Worldpanel data in this work does not imply the endorsement of Kantar Worldpanel in relation to the interpretation or analysis of the data. All errors and omissions remain the responsibility of the author.

[†]Universidad Carlos III de Madrid, Email: acrawfor@eco.uc3m.es

1 Introduction

Storable fast moving consumer goods are frequently the subject of antitrust investigations. For example, recent mergers include Heinz/HP for table sauces, Campina/Friesland for long-life dairy products, Sara Lee/Unilever for personal care products, AB Inbev/SAB Miller in the beer market, and Diageo/Whyte & Mackay in the Scotch whiskey industry.

When assessing antitrust issues in these industries, authorities typically focus on the effect that changes in market structure or firm conduct have on consumer welfare over the course of 1 to 2 years. In practice, as a proxy for changes in consumer welfare, they analyse the likely impact of the merger or alleged firm conduct on prices. Specifically, authorities examine the likelihood that significant non-transitory increases in prices are likely to result from the change in market structure or alleged anti-competitive conduct over the chosen policy horizon.

For storable fast moving consumer goods, authorities often assume that firms compete according to a differentiated Bertrand model. Therefore, the price elasticity matrix that captures substitution patterns implied by consumer responses to permanent price changes over the next 1 to 2 years is a key input into empirical policy analysis.

In antitrust cases, the most commonly used approach is to estimate price elasticities using a static demand model applied to weekly data. However, because static demand models do not incorporate demand dynamics arising from product storability and promotional pricing, they cannot capture inter-temporal substitution.[1] By not acknowledging that short-run volume increases from temporary price cuts draw down on future sales - including those of the promoted product - static demand models provide biased estimates of demand responses to non-transitory price changes.

Specifically, own price elasticities are overstated by static demand models - especially if brand loyalty is prominent. Further, by not including diversion from future sales of rival products, cross-price elasticities tend to be understated. These biases are likely to be most pronounced for the closest substitutes of the promoted good. Therefore, using a static demand model can produce misleading inputs for policy analysis.

When these biased elasticities are used in empirical policy analysis, predicted margins understate market power. Further, because the bias in own and cross price elasticities reinforce one another, diversion ratios are also downward biased - especially for close substitutes. As a result price pressure tests will tend to understate anti-competitive concerns and merger simulation will under-predict price rises.

The source of bias is the mis-specification of the demand model. As highlighted by Erdem et al. (2003), the problem is that inventories are unobserved by researchers. Since current prices and inventories are correlated with past prices, the omission of inventories leads to price endogeneity. As a result, elasticities obtained from a static

[1] In seminal papers, both Erdem et al. (2003) and Hendel and Nevo (2006) show that consumer responses to short-run temporary price changes demand for storable goods are much more elastic than their responses to permanent price changes. It is the latter set of price elasticities that are of interest to antitrust policy makers.

demand estimation are biased. This problem could be resolved if instruments that are correlated with current prices but uncorrelated with past prices were available. However, since observed prices are serially correlated, finding such instruments is challenging - if not impossible.

An alternative is to estimate a dynamic demand model that integrates out over inventories during estimation - a computationally intensive procedure requiring simulation of sequences of purchases, consumption and inventories.[2] Once a dynamic demand model is estimated, the price elasticity matrix measuring demand responses to permanent price changes can be simulated. However, estimating a dynamic demand model is challenging within the timeframe of an antitrust investigation.

In this paper I present an alternative approach to calculating these 'long-run' price elasticities for antitrust policy that can be easily implemented within the timeframe of an antitrust investigation.

This computationally light method produces approximations to price-elasticities adjusted for demand dynamics. The key step is to supplement data on market outcomes used in demand estimation with accounting margins. These margins are often used by firms to assess firm performance and are used as a proxy for market power.[3]

Because profit margins are costly to collect, they are measured over a period of time that exceeds the purchase cycle for storable goods. As such, they contain information on the effect that demand dynamics have on market power. This additional source of information on long-run levels of market power can be used to correct for biased elasticities.

To make use of accounting margins requires a model of the supply side. With demand dynamics and forward looking consumers, firm maximise the present value of profits. The underlying dynamic pricing game is complex - especially many products' prices to set. In addition to optimality conditions, it requires analysis of beliefs of rival firms' pricing behaviour, as well as beliefs about demand responses to price promotions.

The price policy functions that solve the dynamic game describes how each product's price in the next period evolves as a function of past prices and costs. However, it is questionable whether firms possess cognitive and physical resources to solve such a complex game and obtain these policy functions. If not, they need an alternative mechanism to forecast future prices and assess the present value of profits. Instead, given their limited resources, firms are assumed use their accumulated experiences and current market conditions to formulate a statistical model of price forecasts.

Combining the output of a static demand model estimated on weekly data, observed prices, firms' product-level accounting margins, and the price forecasting model, I show how to recover set-valued estimates of reduced-form parameters that measure dynamic diversion ratios.[4]

[2] See Erdem et al. (2003); Hendel and Nevo (2006); Sun (2005); Pires (2016); Wang (2015); Osborne (2017); Crawford (2017) for example of dynamic demand models for storable goods.

[3] See Nevo (2001); Slade (2004); Rojas (2008).

[4] The estimated dynamic diversion set-valued because the both own and cross price elasticities from

The dynamic diversion ratios capture the trade off between the flow of future profits and short-run profit changes in response to short-run price movements when demand is dynamic. They are also a measure of the bias in the price elasticities that results from a mis-specified static demand model. As a result, they can be used to construct set-valued estimates of the price elasticity matrix adjusted for demand dynamics for use in empirical policy analysis.

One potential drawback is that accounting margins imperfectly measure economic profits. However, the use of accounting margins in empirical analysis of mergers is common due to their use in a variety of increasingly common merger screening tools (see Jaffe and Weyl (2013)). In addition to their simplicity, one reason why these methods have become increasingly widespread is the improvement in the quality and availability of detailed cost information.[5] Indeed, notwithstanding the well-documented conceptual differences between the economic and accounting margins, antitrust authorities are placing evidentiary weight on empirical analysis based on firms' margin data.[6]

To illustrate how this method can be employed in practice I apply it to the UK laundry detergent industry between 2002 and 2012. I use it to examine two supply side issues.

First, I explore whether product innovation that reduced the dosage per wash affected the mis-specification bias from static demand models.[7] Shrinking product sizes mean that they take up less storage space and reduce the inventory holding costs. With convex inventory costs, product compaction leads to a relative increase in the demand for large pack-sizes. The larger amount of future consumption purchased in one store visit, the larger the correlation between inventories and prices. As such, one might expect endogeneity bias due to the omission of inventory to increase as a result of the product innovation. In line with this, I find that elasticity biases in the UK laundry detergent industry increase during the phase of intensive product compaction.

Second, I conduct a policy experiment in which I assess whether a cartel in mainland Europe's laundry detergent industry involving the two firms that account for the vast majority of sales in the UK market had any effect in the UK. To test this, I use the 'menu approach' (Bresnahan (1987)) to compare the market power estimates implied by alternative models of competitive interactions during and after the cartel.[8]

I find that without using accounting margins to adjust elasticity estimates, policy

the static demand model are biased. The estimates of the bias are informed by externally measured accounting margins. Therefore, a single data point - the margin - is to identify a minimum of two bias estimates. As a result, the bias estimates are under-identified with two or more products and therefore set-valued.

[5] Antitrust authorities can compel businesses to provide detailed information from management accounts used in the day-to-day running of their commercial activities.

[6] See Pittman et al. (2009) for a recent review of the use of accounting cost information as an input into antitrust policy.

[7] Dosage per wash refers to the physical amount of detergent needed per wash.

[8] Nevo (2001), Hausman and Leonard (2002), Slade (2004), Rojas (2008) and Miller and Weinberg (2017) have assessed market power or merger effects in differentiated product industries using the menu approach.

simulations suggest observed margins are most likely to be produced by anti-competitive conduct.[9] However, in the policy simulations using the set of bias-adjusted price elasticities the case is much less clear-cut. If anything, there is little evidence to support the view that the anti-competitive effects of the European laundry detergent cartel were experienced in the UK.

The remainder of the paper is structured as follows. Section 2 describes the supply side model of price setting when demand is inherently dynamic. Section 3 shows how to combine industry data on market outcomes, accounting margins, a static demand model, and a price forecasting model to recover parameters that adjust price elasticities to reflect consumer responses to permanent price changes when demand dynamics are present. Section 4.1 focuses on the application. First, I describe the UK laundry detergent industry, then estimate a static demand model and apply the methods described in Sections 2 and 3. I use the results to examine the impact of product innovation on the degree of bias in the elasticities produce by the static demand model. Finally, in Section 4.5 I contrast the conclusions of policy simulations using the unadjusted static demand model and the bias-adjusted version. Section 5 concludes.

2 Price setting with demand dynamics

This section presents a model of a firm setting current prices that takes into account expected demand dynamics. In each period firms set prices to maximise the present value of expected profit flows.

When there are no inter-temporal links in demand or costs, the firm's optimisation problem is separable and is solved independently in each time period. However, in many industries, if not most, consumer demand is inherently dynamic. For example, if the firm discounts its product as part of a temporary promotion, consumers may accelerate purchases and stock the product at home for future consumption.

When demand is dynamic, firms consider the impact current prices have on demand today and in the future. The effect of a price change today on expected future profits are reflected in dynamic diversion ratios. These ratios measure the expected future change in profits as a percentage of the contemporaneous profit impact of the price change. They can also be interpreted as approximating the bias from using a mis-specified static demand model to estimate price elasticities.

To evaluate dynamic diversion ratios firms form expectations over the impact of current pricing decisions on future profitability. These expectations depend on demand dynamics and the equilibrium being played in an underlying dynamic pricing game. However, a dynamic demand system is resource intensive to estimate and the underlying promotional

[9] Margins calculated using first order conditions using overly elastic own price elasticities from the static demand estimation understate market power. For the UK laundry detergent in the cartel period, the margins calculated assuming the cartel effected monopoly pricing lie below the observed gross margins. This empirical finding could support a policy makers' case that the cartel on mainland Europe had anti-competitive effects in the UK.

pricing game is a complex frontier research problem - especially when many products are being sold.

The policy functions that solve the underlying dynamic game describe how each product's price in the next period evolves as a function of past prices and costs. However, it is questionable whether firms possess cognitive and physical resources to solve such a complex game and obtain these policy functions. If not, we need an alternative mechanism to forecast future prices.

Acknowledging that solving the game requires a lot of cognitive and physical resources, we assume that, instead, firms take a data driven approach to forecasting prices. Specifically, firms are assumed to use a statistical model of price forecasts based on current market conditions and their accumulated industry experience.

Further, in the baseline version of the dynamic price setting model, it is assumed that nowcasting - contemporaneously updating the price forecasts - may also be a resource intensive activity for firms. As a result, the firm's forecasts are assumed to be taken as given when they make pricing decisions. This simplifies the dynamic pricing problem as current pricing decisions only directly affect future profits through their impact on inter-temporal substitution.

The model is first presented assuming that each firm produces a single product in an oligopoly. Then, it is extended to the multi-product firms. Next, the pricing model is recast in terms of dynamic diversion ratios and I show how they are linked to the bias in the price elasticities estimates that result from a static demand model.

Finally, I show how data on margins, industry outcomes, a static demand model, and estimates of the inter-temporal substitution patterns reflected in dynamic diversion ratios can be used to modify empirical antitrust tools like price pressure tests to account for important demand dynamics.

2.1 Single product oligopoly

Consider an industry with J differentiated products each produced by a different firm. In period t, the price of product j, $p_{j,t}$, and the firm producing it incurs a constant marginal cost of production for each product, $c_{j,t}$.[10] Let $p_t = [p_{1t}, \ldots, p_{Jt}]^\top$ be J-vector recording all prices in period t.

The demand for each product is dynamic. For example, a price promotion in the last τ periods for product j might induce consumers to bring forward purchases to take advantage of the temporarily low sale price. All else equal, this inter-temporal substitution reduces future demand of product j, and - to a lesser extent - current and future demand for substitutes produced by other firms.

To allow for this possibility, the quantity demanded of product j in period t depends on current prices and past prices in the previous τ periods,

$$q_{j,t} = q_j\left(\vec{p}_{t,\tau}\right) \tag{1}$$

[10] It is assumed that fixed costs of production are zero.

where $\vec{p}_{t,\tau} = [p_t, \ldots, p_{t-\tau}]$ denotes the $J \times (\tau + 1)$ matrix containing current prices and the τ-period price history in period t.

Each firm in period t sets the price of its product to maximise the present value of expected profit flows over the next H-periods.[11] The expected NPV of profit for the firms producing product j is

$$\pi_j^{NPV} = \mathbb{E}_t \sum_{h=0}^{H} \delta^h \left(p_{j,t+h} - c_{j,t+h}\right) q_j \left(\vec{p}_{t+h,\tau}\right) \tag{2}$$

where δ is the discount factor.

To evaluate NPV profits each firm form expectations over the future prices of all J products in each of the next H time periods. These expectations depend on firms' beliefs over rivals' future pricing behaviour, their beliefs over consumers' price expectations, as well as their beliefs about the demand response to price promotions. In short, they depend on demand dynamics and the underlying dynamic pricing game.

The underlying dynamic pricing game is complex.[12] Arguably, it is so complex that the cognitive and physical resources needed by firms to solve it in every period are prohibitive. Acknowledging these constraints, it is assumed that firms adopt a boundedly-rational approach and construct a statistical price model that can be estimated using historical market data that records the past prices for all J products.[13] Once estimated, a firm use this statistical model to forecast prices up to H-periods ahead.

Assume that firms use the data on the history of market outcomes to estimate a κ-order Markov process to model the evolution of prices and use it to make price forecasts. That is,

$$p_t = A_0 + \sum_{k=1}^{\kappa} A_k p_{t-k} + u_t \tag{3}$$

where A_0 is a J-vector, each A_k is a $J \times J$ matrix of price coefficients for $k = 1, \ldots, \kappa$ and u_{t+1} is a J-vector of price shocks.

The first order condition used to set the price of product j is

$$\mathbb{E}_t \sum_{h=0}^{H} \delta^h \left(\frac{dp_{j,t+h}}{dp_{j,t}} q_{j,t+h} + \frac{dq_{j,t+h}}{dp_{j,t}} m_{j,t+h} \right) = 0 \tag{4}$$

[11] Where H can be finite or infinite.

[12] The literature characterising equilibrium behaviour in an oligopolistic industries with dynamic pricing is nascent (see Mysliwski et al. (2018)). Mysliwski et al. (2018) examines how firms in the UK butter and margarine industry compete for myopic consumers facing brand switching costs using dynamic promotional pricing strategies. Combining a dynamic demand model with consumer's expectations over future price dynamics with a dynamic price competition model remains a frontier research question.

[13] If the statistical price model used to make forecasts is constructed using data in which the same equilibrium is played, then it can be viewed as reduced firm used by a firm to approximate equilibrium price expectations resulting of the underlying dynamic pricing game.

Noting that the $\frac{\partial p_{k,t}}{\partial p_{j,t}} = 0$ under static Nash Bertrand the first order condition can be re-expressed as

$$q_{j,t} + \frac{\partial q_{j,t}}{\partial p_{j,t}} m_{j,t} + \mathbb{E}_t \sum_{h=1}^{H} \delta^h \frac{dq_{j,t+h}}{dp_{j,t}} m_{j,t+h} + \mathbb{E}_t \sum_{h=1}^{H} \delta^h \frac{dp_{j,t+h}}{dp_{j,t}} q_{j,t+h} = 0 \tag{5}$$

where $q_{j,t}$ is shorthand for the function $q_j(\vec{p}_{t,\tau})$ and $m_{j,t} := p_{j,t} - c_{j,t}$.

The first and the second terms correspond to the terms that enter the first order conditions used by firms in a single product oligopoly to set prices when demand is not dynamic. The additional terms only occur due to the presence of price and demand dynamics.

The third term is the expected change to the discounted net present value of profits earned on sales of product j over the next H-periods due to a change in the current price of product j. It captures the change in future profits that arises due to inter-temporal substitution due to dynamic demand responses. For example, suppose a firm deeply cuts price today to promote product j. Then this term measures the profit lost due to reduction in future sales of product j because consumers accelerate their purchases to take advantage of the sale price.

Finally, the fourth term measures the change in expected revenues from sales of product j over the next H-periods that results from adjustments to future prices in response to changes in the current price of product j. It arises because firm's price forecasts made using κ-order Markov process in equation (3) are adjusted by current prices. Under this forecast model, the dynamic price derivative is a sum of selected parameters in the matrices A_k for $k = 1, \ldots, \kappa$. Noting this, the fourth term can be expressed as a

$$\mathbb{E}_t \sum_{h=1}^{H} \delta^h \frac{dp_{j,t+h}}{dp_{j,t}} q_{j,t+h} = \sum_{h=1}^{H} \delta^h \frac{dp_{j,t+h}}{dp_{j,t}} \bar{q}_{j,t+h} \tag{6}$$

where $\bar{q}_{j,t+h} := \mathbb{E}_t\left[q_{j,t+h}\right]$ is expected quantity of product j sold.[14]

Implicit in above analysis of the fourth term in equation (5) is that assumption that firms' engage in 'nowcasting' - forecasts are updated using prices set in the current period. However, if the process of updating forecasts as a time consuming and resource intensive task, then a boundedly-rational firm might prefer to make price forecasts prior to setting prices in period t.

A1: Firms fix price forecasts prior to setting prices and $\frac{dp_{j,t+h}}{dp_{j,t}} = 0$ for all $j = 1, \ldots, J$ and $h = 1, \ldots, H$

Under assumption A1, the fourth term drops out of equation (5) and the first condition becomes,[15]

[14] Suppose quantity demand of product j in period t is decomposed in to the product of the market size (V_t) and the market share for product j in period t ($s_j(\vec{p}_t)$). That is, $q_{j,t} = V_t s_j(\vec{p}_t)$. Then if the market size is expected to be stable at $\bar{V}$ then $\mathbb{E}_t\left[q_{j,t+h}\right] = \bar{V}\mathbb{E}_t\left[s_j(\vec{p}_{t+h})\right] = \bar{V}\bar{s}_j$ where $\bar{s}_j = \mathbb{E}_t\left[s_j(\vec{p}_{t+h})\right]$ is the expected market share of product j at any given time period.

[15] As discussed below, imposing this assumption considerably simplifies the empirical analysis for the antitrust authority.

$$q_{j,t} + \frac{\partial q_{j,t}}{\partial p_{j,t}} m_{j,t} + \mathbb{E}_t \sum_{h=1}^{H} \delta^h \frac{dq_{j,t+h}}{dp_{j,t}} m_{j,t+h} = 0 \tag{7}$$

In line with perspective that firms face non-negligible time and resource constraints, the baseline model used in the remainder of this paper assumes that firms hold forecasts fixed when setting prices. As such, this fourth term is omitted and equation (7) is used to set p_{jt} .

Remarks Assuming firms used fix price forecasts prior to setting prices yields practical benefits in a policy setting. Specifically, it eliminates the need to estimate an empirical analogue for equation (6) for all products and time periods in the data used in the empirical analysis.

To explore why this is beneficial in a policy setting, suppose the analyst would like to relax this assumption and estimate an empirical analogue for equation (6). This requires the price forecast model in equation (3) to be estimated from market data. However, most industries with dynamic demand contain $J \geq 100$ products in each week. In this case, the forecast model has $J(1 + \kappa J)$ parameters; if $J = 100$ and $\kappa = 1$ then there are over 10,000 parameters to be estimated. With so many parameters the price forecasting model will be challenging to estimate - even in industries with dynamic demand containing only a handful of products.

To circumvent this issue a feasible alternative statistical price forecast model tailored to deal with high-dimensional forecasting problems is required. Moreover, it must maintain the linearity κ-order Markov process in equation (3) that was used to simplify calculation of $\frac{dp_{j,t+h}}{dp_{j,t}}$ in equation (6). This alternative price forecast model, together with any additional assumptions to calculate the empirical analogue to evaluation (6) is discussed at length in section 3.3 below.

In light of the above discussion, the baseline model intended for use in a policy setting assumes that boundedly-rational firms hold fixed price forecasts when setting prices. By retaining the dynamics directly arising from demand dynamics, the resulting model provides the analyst with system of equations that approximates the firms' dynamic pricing problem.

2.2 Multi-product oligopoly

The previous section derived the first order conditions assuming that each of the J products were produced by a separate firms. However, most industries with dynamic demand $N < J$ firms manufacture multiple products in the industry. Next, the system of equations used to set prices is extended to multi-product firms.

Let f index the F firms in the industry and let $\mathcal{J}_f$ denote the set of products sold by firm f. The discounted present value of profits for firm i is,

$$\pi_f^{NPV} = \mathbb{E}_t \sum_{h=0}^{H} \sum_{j \in \mathcal{J}_f} \delta^h \left(p_{j,t+h} - c_{j,t+h}\right) q_j \left(\vec{p}_{t+h,\tau}\right) \tag{8}$$

To extend the analysis to the multi-product case, define $\Delta_{t,\tau}$ as a matrix of demand derivatives for period τ with respect to price changes in period t where $\tau \geq t$ that encodes the ownership structure of the products sold in the industry. If the firm produces both products j and k, the (j,k)-element of $\Delta_{t,\tau}$ is the inter-temporal demand derivative relevant to the pricing of firm i's products,

$$\left[\Delta_{t,\tau}\right]_{(j,k)} = \begin{cases} \frac{dq_{k,\tau}}{dp_{j,t}} & \text{if } j,k \in \mathcal{J}_f \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

The first order condition used to set $p_{j,t}$ for all $j = 1, \ldots, J$ in the baseline model is,

$$q_{j,t} + \sum_{k \in \mathcal{J}_f} \left[\Delta_{t,t}\right]_{(j,k)} m_{k,t} + \mathbb{E}_t \sum_{h=1}^{H} \sum_{k \in \mathcal{J}_f} \delta^h \left[\Delta_{t,t+h}\right]_{(j,k)} m_{k,t+h} = 0 \tag{10}$$

This system of equations serves as the basis for the empirical analysis in the remainder of this paper.

2.3 Dynamic diversion ratios

Before turning to empirical analysis, the system of equations used to set prices (equation (10)) are recast in terms of dynamic diversion ratios.

Dynamic diversion ratios capture the trade-off between the long-run and short run profits when demand is dynamic. In addition to profits lost on foregone sales in the current period, the presence of demand dynamics leads firms to consider the effect of inter-temporal substitution on future profitability. In the context of promotional pricing, the firm weighs the short-run profit gains from a price cut against expected losses due to contemporaneous and inter-temporal cannibalisation. This impact of this can be measured using the dynamic diversion ratios, Ψ_t.

The dynamic diversion ratio between product j promoted in period t and another product k manufactured by firm f is

$$\left[\Psi_t\right]_{(j,k)} := \frac{\mathbb{E}_t \sum_{h=1}^{H} \delta^h \left[\Delta_{t,t+h}\right]_{(j,k)} m_{k,t+h}}{\left[\Delta_{t,t}\right]_{(j,k)} m_{k,t}} \; \forall\, j \in \mathcal{J}_f \tag{11}$$

When multiplied by the change in price for product j, the numerator of the (j,k)-th entry of the Ψ_t matrix is the expected change in future profits from sales of product k over the next H-periods in response to the price change for product j in period t.

Similarly, the denominator is the contemporaneous profit change on sales of product k in the same response to a temporary price cut of product j when multiplied by the price change of product j. Therefore, the (j,k)-th element of Ψ_t measures the ratio of long-run to short-run diverted profits away from product k in response to price cut of product j.

Substituting the expression into equation (10) and rearranging gives

$$q_{j,t} + \sum_{k \in \mathcal{J}_f} [\Delta_{t,t}]_{(j,k)} \, m_{k,t} \left(1 + [\Psi_t]_{(j,k)}\right) = 0 \tag{12}$$

In matrix form, the system of price setting equations is

$$q_t + (\Delta_{t,t} \circ (1 + \Psi_t)) \, m_t = 0 \tag{13}$$

where q_t is a J-vector of quantity demanded in period t, $\Delta_{t,t}$ is a $J \times J$ matrix of static demand derivatives that encodes the ownership of products, Ψ_t is a $J \times J$ matrix dynamic diversion ratios, m_t is a J-vector of mark-ups and $\circ$ denotes the Hadamard product.

Written in this way, demand dynamics in the price setting equations are reflected by the elements of Ψ_t. They capture the degree to which price elasticities are biased due to the omission of demand dynamics. Indeed, without demand dynamics, all elements of Ψ_t are 0 and eq (13) reverts to the static differentiated Bertrand Nash optimality conditions. With demand dynamics, it is the implicit assumption in a static demand model that $\Psi_t = 0$ that leads to biased price elasticities and biased diversion ratios.

To aid empirical analysis, in addition to assuming mark-ups are positive, add two further restrictions on the nature of current and inter-temporal substitution patterns. These are:

A2: Demand for product j is decreasing in own-price in the current period and is increasing in own-price the next H periods. The demand for all other products is such that a cut in the price for product j weakly decreases contemporaneous and future demand over next H-periods for all other goods.

A3: $\left|[\Psi_t]_{(j,j)}\right| \geq [\Psi_t]_{(j,k)}$ $j \neq k,\ j \in \mathcal{J}_f$. Inter-temporal demand responses are strongest with respect to the the products own past prices.

Arguably, these additional restrictions are relatively weak. Restriction A2 requires that purchases of goods in the industries are substitutes both in the period in which they are purchased and in future periods.[16]

Restriction A3 requires that the biases to own-price elasticities are larger in magnitude than biases to cross-price elasticities. This reflects the findings of the existing literature on dynamic demand estimates in storable goods markets.[17] Intuitively, this reflects the fact that inventories built up in response to a promotion have the largest impact on the demand of the product promoted - rather than on the demand of imperfectly

[16]Indeed, antitrust investigations for FMCG industries often focus on markets in which all products are substitutes for one another. As such restriction A2 is consistent with a very large class of industries analysed in a policy setting.

[17]See Erdem et al. (2003); Hendel and Nevo (2006).

substitutable alternatives. Therefore, the biases from omitting inventories - as captured by the parameters in Ψ - are largest for own-price elasticities.

Figure 1 provides a graphical representation of the restrictions that A2 and A3 place on inter-temporal substitution patterns. The left panel shows the percentage change in sales of product j when the price of product j is temporarily cut by 1% in period 1. The right panel shows the corresponding percentage change in sales of product k.

The figure shows that initially, the sales of product j rise in response to the temporary price cut. Also, because product k is a substitute for product j, the temporary price cut of product j diverts sales away from product k in the first period.

In the subsequent periods, both product j and product k's sales are lower than those without the price cut - albeit at a diminishing rate. The loss of future sales reflects that, in part, some of the demand uplift in period 1 is drawn from future demand. That is, future sales of products j and k are substitutes for current sales of product j. However, the percentage change in future sales of product k are less pronounced than product j. This is line with restriction A2; consumers consider future sales of product k to be a more distant substitute for product j in the current period than future sales of product j.

This figure also highlights why estimates of medium to long-run price elasticities using a mis-specified static demand model are biased. Implicitly, a static demand model omits the change in sales from period 2 onwards. As a result, future reductions in sales of product j are not netted off from the immediate demand response to the price cut. Therefore the own-price elasticity of product j is overstated. Likewise, by not taking into account the loss of future sales of product k in response to temporary price cut of product j, the cross-price elasticity of k with respect to j is understated.

Adding these restrictions to profit maximising behaviour places bounds on the elements of Ψ_t. These bounds are shown in Proposition 1. It states that the diagonal elements of Ψ lie in an interval between -1 and 0 and off diagonal terms are bounded below at 0 and above by the absolute value of the diagonal element on the same row.

Proposition 1: Suppose A1-A3 hold. The prices in each period are set using the following system of equations,

$$q_t + \left(\Delta_{[t,t]} \circ (1 + \Psi_t)\right) m_t = 0 \tag{14}$$

$$0 \leq 1 + [\Psi_t]_{(j,j)} \leq 1 \qquad \forall\, j \in \mathcal{J}_f \tag{15}$$

$$1 \leq 1 + [\Psi_t]_{(j,k)} \leq 1 - [\Psi_t]_{(j,j)} \leq 2 \quad \forall\, j \neq k, j \in \mathcal{J}_f \tag{16}$$

where $m_t := p_t - c_t$ and $\circ$ denotes the Hadamard product.[18]

Before turning to empirical analysis, I describe how the system of price settings equations (14-16) can be used in existing antitrust policy tools often applied to industries in which demand dynamics are common.

[18]See Annex A.1 for the proof of Proposition 1.

Figure 1: Inter-temporal substitutions patterns of product j and k over 10 periods in response to a temporary 1% price cut of product j at $t = 1$

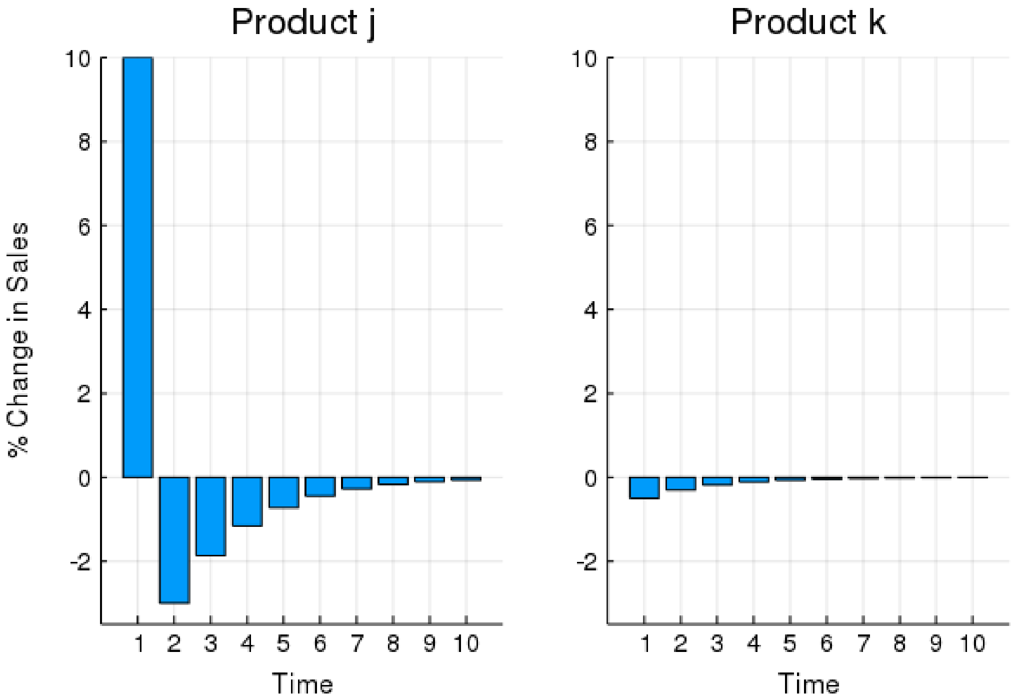

2.4 Antitrust policy: price pressure tests

Next I how price pressure tests can be modified using Ψ_t to (approximately) correct for these biases that result from omitting the industry demand dynamics.

The 2010 US Horizontal Merger Guidelines introduced upwards price pressure measures UPP (Farrell and Shapiro (2010)) and the gross upwards price pressure index (GUPPI) (Moresi (2010)). Since then, these price pressure tests have been used to screen mergers, aid market definition and assess closeness of competition - especially in FMCG industries. This section shows how dynamic diversion ratios can be used to adjust diversion ratios used in upward price pressure tests to reflect demand dynamics in GUPPI formula.

To understand how Ψ is linked to GUPPI, first define the ratio of short-run profits earned on sales of product k diverted away from short-run profits earned on sales of product j when the price of product j changes,

$$DR^{\pi}_{j,k} := -\frac{[\Delta_{t,t}]_{(j,k)}}{[\Delta_{t,t}]_{(j,j)}} \frac{(p_{k,t} - c_{k,t})}{(p_{j,t} - c_{j,t})} \tag{17}$$

Recall, that GUPPI is defined as follows

$$GUPPI_{j,k} := -\frac{[\Delta_{t,t}]_{(j,k)}}{[\Delta_{t,t}]_{(j,j)}} \frac{(p_{k,t} - c_{k,t})}{p_{j,t}} = DR^{\pi}_{j,k}\mu_{j,t} \tag{18}$$

where $\mu_{j,t}$ is the percentage mark-up on product j in period t. Then define a dynamic counterpart of $DR^{\pi}_{j,k}$ as the the ratio of expected NPV of profits earned on sales of product k to profits lost on sales of product j over the next H-periods as a result of a price change of product j in period t,

$$DDR^{\pi}_{jk} := -\frac{\mathbb{E}_t \sum_{h=0}^{H} \delta^h \left[\Delta_{t,t+h}\right]_{(j,k)} \left(p_{k,t+h} - c_{k,t+h}\right)}{\mathbb{E}_t \sum_{h=0}^{H} \delta^h \left[\Delta_{t,t+h}\right]_{(j,j)} \left(p_{j,t+h} - c_{j,t+h}\right)} = \frac{1 + \left[\Psi_t\right]_{(j,k)}}{1 + \left[\Psi_t\right]_{(j,j)}} DR^{\pi}_{j,k} \tag{19}$$

Replacing the $DR^{\pi}_{j,k}$ with its dynamic counterpart, $DDR^{\pi}_{(j,k)}$ defines a GUPPI adjusted for inter-temporal substitution arising from demand dynamics,

$$DGUPPI_{j,k} = DDR^{\pi}_{j,k}\mu_{j,t} = \frac{1 + \left[\Psi_t\right]_{(j,k)}}{1 + \left[\Psi_t\right]_{(j,j)}} GUPPI_{j,k} \tag{20}$$

To illustrate how the dynamics-adjusted GUPPI can be used to capture dynamics once the elements of $\bar{\Psi}$ have been estimated, consider a firm assessing the impact on profits of its promotional pricing decisions on two substitutable products j and k over a one month planning horizon.

Suppose the firm producing product j is considering temporarily cutting its price from £15 to £10 for the current week. If it does so, it expects to generate an additional 100 sales as a result. However, the firm expects that 20 of those sales would occur at the regular higher price of £15 within the subsequent week over the next month without the price cut.

Assuming a marginal cost of £5 in all periods, 40% of short-run profits made on sales of good j immediately following the price cut come are drawn from expected profits from future profit flows. That is, the (j, j)-th element of $\bar{\Psi}$ is -0.4.

Further suppose, the effect of the price cut of product j on product k is the loss of 10 sales today and a further 3 sales over the next month. Assuming its price and cost remain unchanged over the period and its sales yield a £10 margin, the firm anticipates that consumer demand dynamics will lead to an additional 30% of losses on top of short-run profits lost on product k over the next month. Equivalently the (j, k)-th entry of $\bar{\Psi}$ is 0.3.

In this simple example, only 20% of the firm's profits are the result of sales diverted away from product k in the period of the price change and the GUPPI from k to j is 0.1. Recall, that the (j, j)-th and the (j, k)-th entry of $\bar{\Psi}$ are -0.4 and 0.3, respectively. Therefore, the GUPPI after anticipated demand dynamics are factored in is 0.217 - more than double the GUPPI calibrated with contemporaneous profit movements alone. This higher dynamic diversion ratio demonstrates that products j and k are much closer substitutes for non-transitory price changes than implied by a static analysis.

3 Identification and Estimation

Treating the dynamic diversion ratios as reduced form parameters, I show in this section how they can be identified and estimated by combining accounting margins with data on industry outcomes.

In general, the number of parameters in the dynamic diversion ratios that enter the price setting equations over the interval in which the accounting margins are recorded exceed the number of accounting margins available to the analyst. As a result they are under-identified and therefore set-valued.

First, the data assumed to be available to the analyst is described. Next, the additional assumptions imposed to give the set-identified parameters of Ψ_t enough empirical content for use in a policy setting are discussed. The section is closed by outlining the additional steps needed to extend the empirical approach to relax assumption A1 and allow firms to update price forecasts in the period in which they and their rivals set prices. However, as noted above, a τ-order Markov process is unlikely to be estimable in many FMCG industries due to the number of products sold. Therefore, this final section also describes a factor model that firms are assumed to use for price forecasting.

3.1 Available data

Throughout this section I assume that the data set has two components: (i) data on accounting margins, and (ii) data on industry outcomes.

Accounting Margins Accounting margins are increasingly widely used in antitrust policy as a proxy for market power. When measured over a relatively long time horizon (i.e. a quarter or year) they capture the average effect of inter-temporal substitution on market power. This feature of margins is used to set-identify values for the elements of Ψ_t.

In practice margins are costly to collect and collate. Therefore firms often report accounting margins aggregated over some partition of its product space (i.e. brands or size) and report them on a periodic basis (i.e. annually). For example, the empirical application in this paper uses annual firm-wide gross margins taken from annual reporting data. However, in an antitrust investigation, it is possible that through access firms' management accounts a more disaggregated dataset of margins can be collected. The remainder of this section therefore allows for the possibility that the analyst has access to margin data for N groups of subsets of products measured over an interval of $[1,\bar{T}]$ periods. Let J_n denote the number of products in group $n = 1, \ldots, N$.

Market data The data on industry outcomes is either aggregated market data or consumer level micro-data. In both cases, it is used to estimate a static demand model. Once estimated, the demand model is combined with industry data on prices and products sold to construct purchase probabilities and demand derivatives in each period in which the data is available.

It is assumed that the data on industry outcomes is available over the same interval that the accounting margins are measured over. If so, the output from the static demand model can be used to populate the system of price setting equations (eqs. (14), (15), and (16)) to estimate sets of values for the reduced form parameters in the dynamic diversion ratios, $\{\Psi_t\}$ for $t = 1, \ldots, \bar{T}$.

For example, the empirical application in this paper uses consumer level micro-data in each week from 2002 to 2012 to estimate a demand model together with yearly gross margins taken from annual reports over 2002 to 2012 for the two large companies that dominate the industry in the application.

3.2 Identification of Ψ

The number of reduced form parameters in $\{\Psi_t\}$ for $t = 1, \ldots, \bar{T}$ in the system of price setting equation (14), (15), and (16)) for products in group n is $\bar{T}J_n^2$. To aid identification of Ψ_t we impose further restriction on dynamic diversion ratio that reduce the number of parameters to be estimated to J_n^2.

A4: The expected value of numerator in the dynamic diversion ratios is covariance stationary. Therefore

$$\Psi_t = \bar{\Psi} \,\forall\, t = 1, \ldots, \bar{T} \tag{21}$$

Under assumption A4, the system of price setting equations for products in group n for all $t = 1, \ldots, \bar{T}$ over which the external accounting margin is measured is

$$
\begin{aligned}
q_t + \left(\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}\right)\right) m_t = 0 & & (22)\\
0 \leq 1 + \left[\bar{\Psi}\right]_{(j,j)} \leq 1 & \qquad \forall\, j \in \mathcal{J}_n & (23)\\
1 \leq 1 + \left[\bar{\Psi}\right]_{(j,k)} \leq 1 - \left[\bar{\Psi}\right]_{(j,j)} & \quad \forall\, j \neq k, j, k \in \mathcal{J}_n & (24)
\end{aligned}
$$

Rearranging the system of price setting equations for J_n products in group n (eq. (22)), the markup in each period can be expressed as a function of demand short-run demand derivatives and dynamic diversion ratios

$$p_t - c_t = -\left(\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}\right)\right)^{-1} q_t \,\forall\, t = 1, \ldots, \bar{T} \tag{25}$$

where $\left[\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}\right)\right]$ is an invertible matrix of dynamically adjusted demand derivatives.[19] Taking the average over $\bar{T}$ periods of the quantity-weighted average mark-up of products in group n gives an equation for the gross margin, $\mu_{n,[1:\bar{T}]}$.

[19]The non-singularity of this matrix is guaranteed if its columns are linearly independent. As shown in Annex A.4, this requires that the aggregate long-run diversion ratios for products in group n produced by firm f differ. This is arguably a weak requirement given that profit maximising FMCG firms are unlikely choose a portfolio of differentiated products with symmetric substitution patterns .

Proposition 2: Product group n's margin measured over $t = 1, \ldots, \bar{T}$ periods can be expressed as a function of the sequence of prices, purchase quantities, short-run demand derivatives and dynamic diversion ratios,

$$\mu_{n,[1:\bar{T}]} \quad = \quad -\frac{\sum_{t=1}^{\bar{T}} q_{n,t}^{\top} \left(\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}\right)\right)^{-1} q_t}{\sum_{t=1}^{\bar{T}} q_{n,t}^{\top} p_t} \tag{26}$$

where $q_{n,t}$ is vector of quantities purchase whose j-th entry is $q_{j,t}$ if $j \in \mathcal{J}_n$ and 0 otherwise.[20]

provides the analyst with an equation in which the only unknowns are the J_n^2 elements of $\bar{\Psi}$.

Sets of values of the J_n^2 elements of $\bar{\Psi}$ that can be estimated by plugging in the observed accounting margin for group n of products measured the interval $[1, \bar{T}]$ into equation (22) together with observed prices, estimated purchase shares and demand derivatives from the static demand model for each period $t = 1, \ldots, \bar{T}$.

However, with only 1 external margin in equation (26) together with $2J_n$ inequality constraints (eq. (23) and (24)), the J_n^2 parameters in $\bar{\Psi}$ are not point-identified.[21] Instead, the solution to this system of equations defines sets of values that the elements in $\bar{\Psi}$ can take.

The remainder of this section sets out additional assumptions used to reduce the number of parameters in $\bar{\Psi}$ to be estimated for products in one of the groups, n. One option is to analyse the nature of the data collected to see if economically plausible additional restrictions can be added to reduce the number of parameters to be estimated. To that end, I add two further assumptions.

First, without variation of product margins within the n groups of products the available data does allow the analyst to draw inference about individual $\left[\bar{\Psi}\right]_{(j,j)}$ for each the j products in group n. As such, I introduce a further assumption that constrains all the dynamic diversion ratios on the diagonal of $\bar{\Psi}$ for all products in group n to be the same.

A5: $\left[\bar{\Psi}\right]_{(j,j)} = \psi_n^{own}$ for all $j \in \mathcal{J}_n$ and $n = 1, \ldots, N$.

Second, because the cross-section of observed margins has no additional information on covariances between mark-ups and demand derivatives, the off-diagonal elements are not identified. Therefore, I approximate the off-diagonal elements of matrix of dynamic diversion ratios by assuming that they are the same for any two products, $k \neq j$ and $k' \neq j$.

A6: $\left[\bar{\Psi}\right]_{(j,k)} = \psi_n^{cross}$ for all $j \neq k, j \in \mathcal{J}_n$ and $n = 1, \ldots, N$

Adding assumptions A5 and A6 further reduces the number of parameters from J_n^2 to 2 parameters for each group n. With only these 2 parameters the analyst uses equation

[20] See Annex A.3 for a proof of Proposition 2.

[21] For the non-trivial case of $J_n > 1$.

(26) to trace out sets of parameters $(\psi_n^{own}, \psi_n^{cross})$ that also satisfy the modified bounds

$$0 \leq 1 + \psi_n^{own} \leq 1 \tag{27}$$
$$1 \leq 1 + \psi_n^{cross} \leq 1 - \psi_n^{own} \tag{28}$$

3.3 Price forecasting

This section discusses how to extend the estimation of $\bar{\Psi}$ and relax assumption A1. In the description of the model in section 2.1, it was assumed that firms use the data on the history of market outcomes to estimate a κ-order Markov process to model the evolution of prices and use it to make price forecasts.

$$p_t = A_0 + \sum_{k=1}^{\kappa} A_k p_{t-k} + u_t \tag{29}$$

where A_0 is a J-vector, each A_k is a $J \times J$ matrix of price coefficients for $k = 1, \ldots, \kappa$ and u_{t+1} is a J-vector of price shocks.

However, a noted above, most industries with dynamic demand contain $J \geq 100$ products in each week. In this case, the forecast model has $J(1 + \kappa J)$ parameters; if $J = 100$ and $\kappa = 1$ then there are over 10,000 parameters to be estimated. With so many parameters the price forecasting model will be challenging to estimate - even in industries with dynamic demand containing only a handful of products.

Moreover, tracking and forecasting $J \geq 100$ product's prices is a cognitively challenging and resource intensive task for firms. Indeed, boundedly-rational firms may optimally choose a more parsimonious forecasting model.

Instead, it is assumed that, through their repeated experiences of selling products in the industry firms are assumed to be able to relate individual products price movements to a low dimensional set of price trends that underpin observed price dynamics. If so, firms need only to track and forecast this low dimensional set of price trends. In line with this bounded-rationality approach to forecasting prices, an alternative forecast model is

$$p_{j,t} = \lambda_j^\top f_t + \epsilon_{jt} \tag{30}$$
$$f_{t+1} = \sum_{k=0}^{\tau} A_k f_{t-k} + u_{t+1} \tag{31}$$
$$\begin{bmatrix} \epsilon_t \\ u_{t+1} \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \Sigma_{pp} & 0 \\ 0 & \Sigma_{ff} \end{bmatrix} \right) \tag{32}$$

where f_t is an R-vector of low dimensional set of price trends or 'factors'. The relationship of individual product to these price factors is captured by a R-vector of product specific weights λ_j or 'factor loadings'.

Using this forecasting model the $J \times J$ matrix of expected changes to forecast prices h-periods ahead as a result of price changes in the current period is given by[22]

$$\mathbb{E}_t\left[\frac{dp_{t+h}}{dp_t}\right] = \Lambda A^h f_t f_t^\top \Lambda^\top \left[\Lambda f_t f_t^\top \Lambda^\top + \Sigma_{pp}\right]^{-1}$$

where $\Lambda = [\lambda_1, \ldots, \lambda_J]$ is a $J \times R$ matrix of time-invariant factor loadings. Noting that this is a $J \times J$ matrix of estimated parameters, we can write the multi-product equivalent of the price forecasting terms omitted under assumption A1 as

$$\mathbb{E}_t \sum_{h=1}^{H} \sum_{k \in \mathcal{J}_n} \delta^h \frac{dp_{k,t+h}}{dp_{j,t}} q_{k,t+h} = \sum_{k \in \mathcal{J}_n} [\Gamma_t]_{(j,k)}\, \bar{q}_k \tag{33}$$

where

$$\Gamma_t := \sum_{h=1}^{H} \delta^h \Lambda A^h f_t f_t^\top \Lambda^\top \left[\Lambda f_t f_t^\top \Lambda^\top + \Sigma_{pp}\right]^{-1} \tag{34}$$

is a $J \times J$ matrix of impulse response for the forecast price of product k in response to change in product j in the current period t and $\bar{q}$ is a J-vector of expected product sales. It is assumed that $\mathbb{E}_t[q_{k,t+h}] = \bar{q}_k = \bar{T}^{-1} \sum_{t=1}^{\bar{T}} q_{j,t}$ where $q_{j.t}$ is calculated using the static demand model and industry data.

Using this notation above, system of price setting equations when assumption A1 is relaxed for a products in group n in period t is

$$q_t + \left(\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}\right)\right) m_t + \Gamma_t \bar{q} = 0 \tag{35}$$
$$0 \leq 1 + \psi_n^{own} \leq 1 \tag{36}$$
$$1 \leq 1 + \psi_n^{cross} \leq 1 - \psi_n^{own} \tag{37}$$

and the corresponding estimation equation is

$$\mu_{n,[1:\bar{T}]} = -\frac{\sum_{t=1}^{\bar{T}} q_{n,t}^\top \left(\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}\right)\right)^{-1} \left(q_t + \Gamma_t \bar{q}\right)}{\sum_{t=1}^{\bar{T}} q_{n,t}^\top p_t} \tag{38}$$

Assuming the demand model has been estimated, estimation of equation (38) has two steps.

The first is to estimate the price forecasting model using price data from market outcomes. The model can be estimated using factor a model approach imposing scale and rotation normalisations for factors and loadings.[23] Statistical criteria can be used

[22]See Annex A.5 for a derivation

[23]see Bai and Ng (2008) for a review of factor models.

to select the number of factors R and number of lags in the factor VAR, τ. Industry details can be used to choose H.

Allied with estimates of $\bar{q}$ from the demand model, empirical estimates of the sequence of additional terms $\Gamma_t \bar{q}$ for $t = 1, \ldots, \bar{T}$ can be constructed using estimated factors, loadings, VAR coefficients, and the price covariance matrix Σ_{pp}. Next, the analyst solves for the the set of values for $\{\psi_n^{own}, \psi_n^{cross}\}$ in $\bar{\Psi}$ that satisfy equations (36), (37) and (38).

4 Application: UK laundry detergent

In this section, I apply the model described in Sections 2 and 3 to the UK laundry detergent sector. I begin by providing an overview of the industry. In particular, I highlight how technological innovations have shrunk products, reduced storage costs and led to changes in pricing behaviour. Intuitively, reductions in product size enable households to hold higher levels of inventory without incurring additional costs.

Since the source of bias in the static demand estimation is due to the omission of inventories, both changes in storage costs and intensification of price competition due to product compaction increase inter-temporal substitution. In particular forward-looking consumers become ever more responsive to short-run price cuts. Because a static demand model cannot capture the offsetting reduction in future demand the due to inter-temporal factors, it is likely that the degree of bias is exacerbated.

Next, I analyse whether there is any evidence to support this hypothesis. To that end, I first specify a static demand model and estimate it using household purchase diary data from Kantar for a large UK supermarket from 2002 to 2012. Then using publicly observed margins in each year for firm A and firm B, using the methodology described in Section 3 I recover the sets of dynamic diversion ratios in each year. Combining the demand model with sets of estimates for $\bar{\Psi}$, I show how product innovation that lowers storage costs affects estimates of price elasticities. I find evidence to support the hypothesis that bias increases due to product compaction. In turn, highlighting how this method can be effectively use to correct for biases due to omitted demand dynamics.

Finally, I use the estimates to test whether there is evidence that a laundry detergent cartel in mainland Europe between 2002 and 2005 affected the UK detergent market. I find that using unadjusted elasticity estimates in policy simulations suggest observed margins may have been the result of anti-competitive conduct. However, using bias-adjusted price elasticities the analysis is much less clear-cut. In this case, there is little evidence of anti-competitive effects in the UK. This exercise highlights that the omission of demand dynamics has the potential to lead to misguided policy conclusions.

4.1 The UK laundry detergent industry

This section describes the UK laundry detergent industry. It provides an overview of the type of detergent products and brands sold at a leading UK retailer. Subsequently it compares how industry outcomes have been shaped by product innovation and how pricing strategies have evolved. Where applicable I discuss how these industry trends might impact on the estimation of price elasticities using mis-specified static demand models. First, I describe the data used in this section and the remainder of the paper.

4.1.1 Data

The analysis of the UK laundry detergent industry is based on individual household purchase data from 1st January 2002 until 31st October 2012.

Households that take part in the survey scan the barcode of the items they purchase. Using the scanned barcode, the survey records the price and number of packs bought together with the characteristics of the product purchased. In addition, the purchase date and store in which the product was bought is also recorded. The purchase data is supplemented by annually updated household demographics. These include data on composition of the household, social class status, and sundry features of the household.

In the remainder of the paper, to avoid complexities related to store choice, the analysis is conditioned on detergent purchased from a single leading UK's supermarket. Prices charged by supermarkets the same fascia are the same stock keeping unit (SKU) in different stores in excess of 280 sq. ft in size.[24] As such, purchases across different stores within a particular supermarket fascia in the UK can be pooled for the purposes of empirical analysis.

4.1.2 Industry Overview

The UK laundry detergent industry is populated by a diverse array of brands, formats, and pack sizes. They are sold in SKUs, each containing a single type of detergent. In general, a detergent is defined by its format, brand and the chemical properties of the enzymes it contains (i.e. non-bio/bio, stain removal properties, scent etc). In the remainder of this paper I focus my analysis on the major distinguishing features of a SKU - its brand and format.

The UK laundry detergent industry is dominated by firm A and firm B. Table 1 shows that together they account for around 75 to 85 percent of households' annual purchases of laundry detergents. Outside of these two major producers of branded products, the retailer's private label products commands the largest share, although its share has declined from 25 percent in 2002 to 13 percent in 2012. A fringe of small niche brands account for the remainder of SKUs sold.

[24] This pricing policy follows an undertaking in 2001 following a market investigation by the UK Competition Commission into the groceries industry in 2000.

Table 1: Firm shares of expenditure

Firm	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012
Firm A	43.8	46.0	44.4	51.2	51.6	50.1	48.6	47.4	53.9	52.0	50.5
Firm B	29.7	29.8	31.5	29.1	28.6	30.6	32.8	36.3	31.9	32.3	34.1
Private Label	24.7	22.3	21.5	16.9	17.4	16.8	16.3	14.3	12.9	13.8	12.4
Other	1.8	1.9	2.5	2.8	2.5	2.5	2.3	2.0	1.4	1.9	3.0

Source: Kantar WorldPanel

Laundry detergent is sold in six formats: powder, liquid, tablets, liquid capsules, super concentrated liquid, and gel. Figure 2 shows how each format's share of household spend evolved from 2002 to 2012. Powder and tablets, the most popular formats in 2002, saw a notable decline in their market share from 2006. Initially, the market share ceded by tablets and powder products was largely captured by the new super concentrated liquid laundry detergents. Subsequent declines in market share, especially for tablets, coincide with the launch of gel products in 2008. Following their introduction both super concentrated liquids and gel products quickly gained market share; by 2010 they were the second and third most purchased format respectively. Liquid capsules steadily accumulated market share from 8 percent to 15 percent over the sample period.

Figure 2: Share of SKU purchases by format from 2002 to 2012

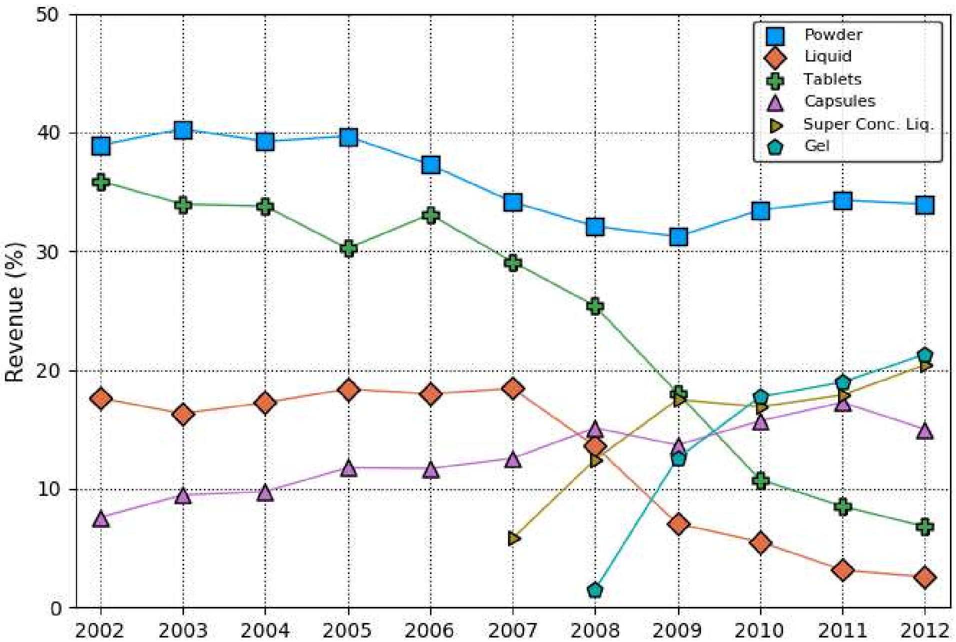

In addition to the wide variety of formats, laundry detergents can be purchased in many different pack sizes.[25]

Because formats are defined in different dosage metrics, I use the number of washes in the pack as a common metric across all SKUs. Not only do dosage metrics vary across formats, they also evolve over time within formats. Initially, this was due to a series of industry-wide initiatives that sought to reduce the environmental impact of the production and use of laundry detergents. Later, further product compactions were the result of firm specific innovations.

For most formats, these initiatives served to decrease the dosage per wash over time. A corollary of the reduction in dosage per wash is that more washes can be included in smaller pack sizes. The impact of the evolution in the dosage per wash and the number of washes per pack is shown in Figure 3.

Most notably, for powder products: from 2002 to 2012 the average dosage per wash fell by over 20 percent. Over the same period, the average powder SKU nearly doubled in size - from 18.5 to 34.5 washes. Further, the average super concentrated liquid SKU in 2012 contains 35 washes – which is roughly twice as many washes as the corresponding regular liquid product by 2012 with 17 washes.

The dosage per wash reduction lowers households' cost of storage per wash and exerts downward pressure on firms' transport and packaging costs per wash sold. Figure 4 shows that, as expected, these cost savings have contributed to the increasing popularity of larger packs of laundry detergent.[26] In 2002, around 75% of household spend was on SKUs with fewer than 24 washes, 10 years later this figure was less than 35 percent.

[25]The UK laundry detergent industry differs in this respect from the one studied by Hendel and Nevo (2006). They restrict attention to powder products and examine brand choice conditional on size choice from a small number of discrete sizes: 16oz, 32oz, 64oz, 96oz, and 128oz. Erdem et al. (2003) also focus on only five different weight choices in the US Ketchup market in their dynamic demand estimation.

[26]To maintain a consistent measure of size, SKUs are grouped by the number of washes they contain. Partitioning the sample into groups whose boundaries are defined by 25th, 50th and 75th percentiles of purchased washes from 2002 to 2012 results in groups with 0-17 washes, 17 to 24 washes, 25 to 40 washes and 41 or more washes respectively.

Figure 3: Compaction and concentration of detergent: Dosage and SKU size from 2002 to 2012

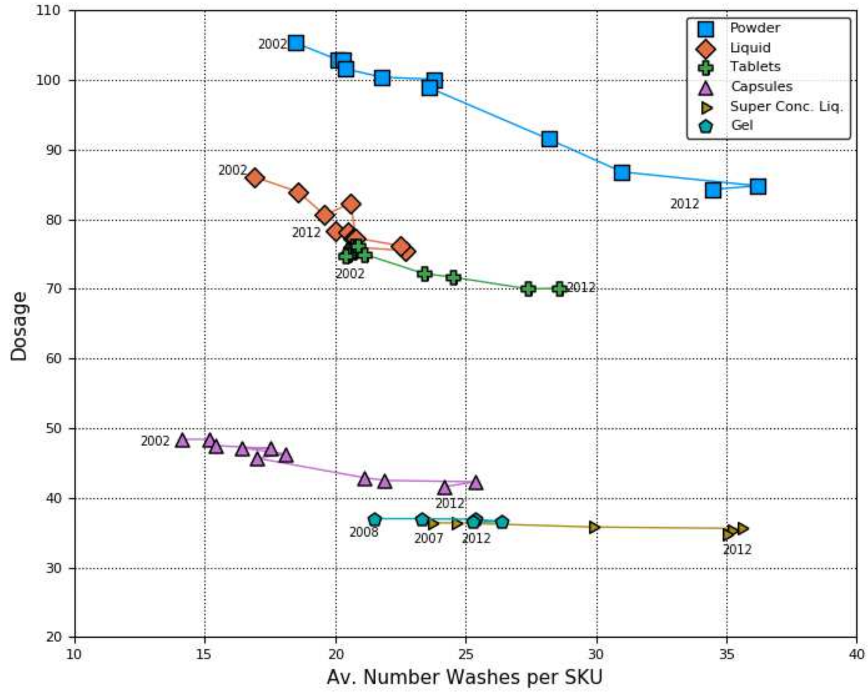

Figure 4: Number of washes: share of SKUs purchased 2002 to 2012

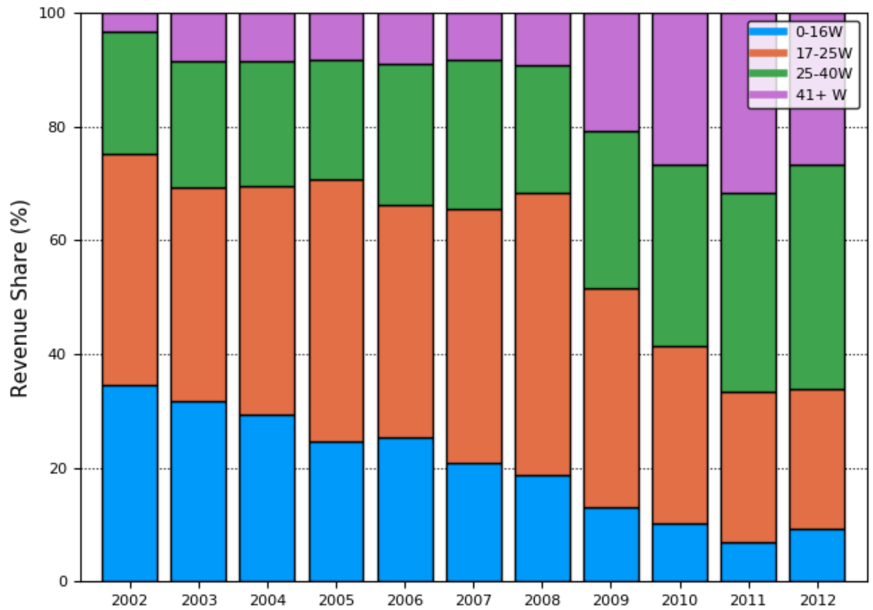

Product compaction also had an impact on pricing behaviour. Figure 5 uses a series of box plots to display the distribution of prices per wash in each quarter from 2002 to

2012. The top panel shows the price per wash distribution for firm A and bottom panel shows this distribution for firm B.

For both firm A and firm B, the whiskers and inter-quartile range of box plots from first quarter in 2002 up to the final quarter in 2006 are relatively constant, if anything they narrow after 2004. Over this period, the whiskers tend to lie between 10p and 30p per wash and the interquartile ranges lie between 15p and 23p per wash. From the first quarter in 2007 onwards, the whiskers and inter-quartile ranges of the box plot fan out for both firm A and firm B.

This increased price dispersion coincides with the introduction of new products and the ending of a detergent cartel in mainland Europe that was effect from 2002 to 2005 (see Section 4.5). Increased price dispersion has been found in other industries following the removal of pricing restrictions.[27] As such, it is possible that increased price dispersion reflects the restoration of competitive forces following the removal of anti-competitive pricing and advertising restraints imposed by a cartel. As such, these price movements may reflect efforts from firm A and firm B to increase the degree of differentiation of the brands from 2006 onwards.

However, there may be alternative explanations for the observed change in pricing strategy. For example, pricing strategy might change in response to increased consumer price sensitivity (i.e. perhaps due to a perceived drop in income after the financial crisis) or new format adoption as new compacted product are rolled out.

The red line in Figure 5 plots the average posted price per wash and the green line plots the average price per wash of purchased products. It shows that the average posted price per wash is steadily increasing for firm A, but broadly constant for firm B. However, the average purchased price per wash diverges from the posted price per wash, and this gap increases over time. This is especially pronounced for firm B, with divergence starting as early as the second quarter in 2005. From 2010 onwards, the average purchased price tracks close to the lower quartile of the posted price distribution.

On one hand, these pricing patterns might reflect the change in the pricing behaviour of UK laundry detergent suppliers in the period that followed the ending of the cartel in 2005. On the other, it may reflect the impact of product compaction. A corollary of product compaction is that households can purchase a higher fraction of their detergent demand in a single purchase. In turn, this may have increased the intensity of competition and contributing to deeper and more frequent discounting. Other explanations include increased price sensitivity of households due to changes in the macroeconomic climate, or changes in pricing strategy to encourage new product adoption.

Like product compaction, the increasingly widespread use of promotional pricing strategies has the potential to exacerbate biases in estimates of own-price elasticities by static demand models. By increasing the opportunity cost of mis-timing purchases,

[27] Frank and Salkever's (1997) study of post-liberalisation pricing of pharmaceutical products finds that prices of premium branded products typically increase whereas the price of standard branded products tended to fall.

Figure 5: Price per wash (PPW) distribution: firms A and B from 2002 to 2012.

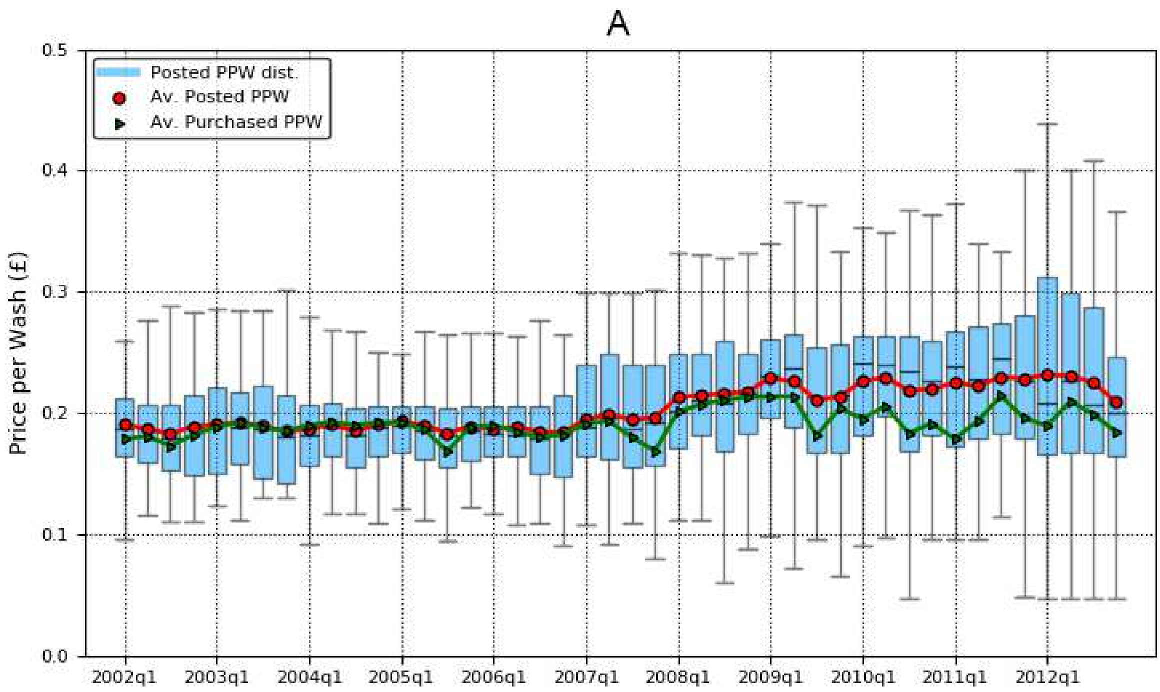

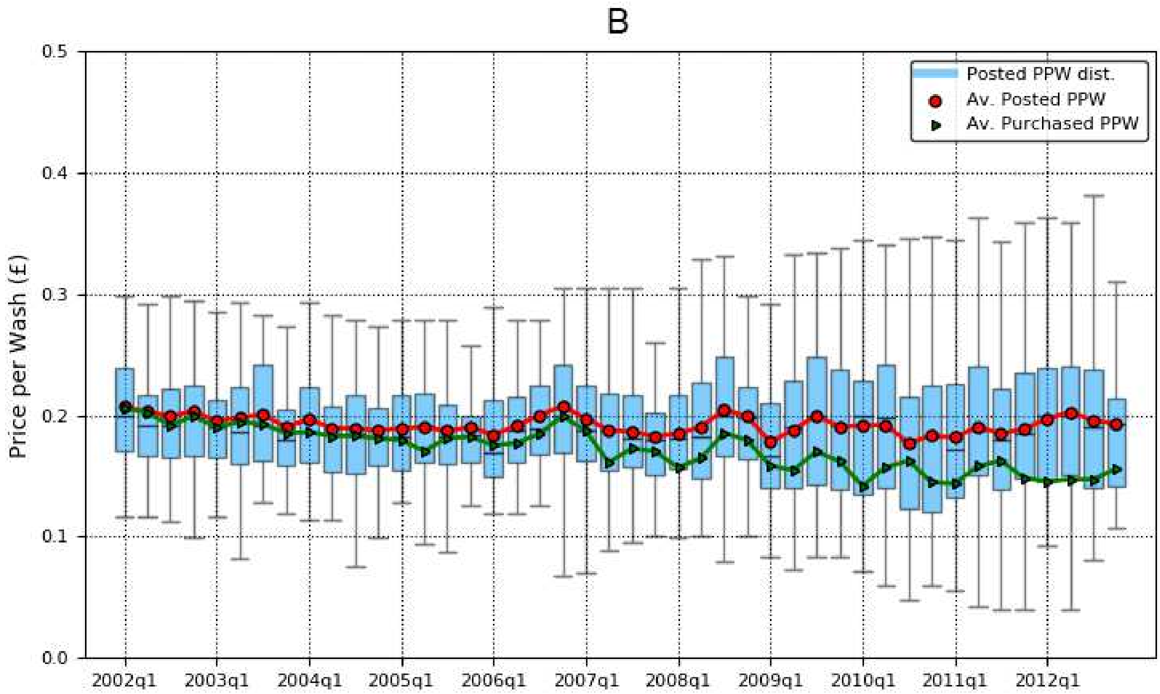

consumers have an increased incentive to stock up during promotions. Not being able to capture the more pronounced inter-temporal substitution this creates, static demand models reflect this heightened short-run price sensitivity by overstating (understating) own-price (cross-price) elasticities.

4.2 Demand for laundry detergent

In this section I set out a static demand model of the UK laundry detergent and estimate it using the purchase diary data from Kantar data described in Section 4.1.1. I estimate two demand models that are commonly used in antitrust policy; conditional logit and nested logit.

4.2.1 Demand model

Detergents are sold in one of $j = 1, \ldots, J$ stock keeping units (SKUs). A SKU is defined by the detergent it contains and the number of washes it contains. The SKUs sold in market t, $\mathcal{J}_t$, are manufactured by $f = 1, \ldots, F$ firms. The outside good is denoted by $j = 0$ and represents the decision not to purchase in market t.

Household i elects to purchase good $j^\star$ from a market t to maximise conditional indirect utility, $V_{ijt}(x_j, p_{jt}, z_{it}; \theta)$

$$j^\star = \arg\max_{j \in \mathcal{J}_t} V_{ijt}(x_j, p_{jt}, z_{it}; \theta) \tag{39}$$

where p_{jt} is the price of SKU j in market t, x_j is a K-vector of SKU attributes, z_{it} is an L-vector of household characteristics, and θ is the set of parameters entering the conditional indirect utility function.

The conditional indirect utility for household i from purchasing product j in market t is

$$V_{ijt}(x_j, p_{jt}, z_{it}; \theta) = x_j^\top \beta_{it} + \alpha_{it} p_{jt} + \epsilon_{ijt} \tag{40}$$

The observed product attributes include the number of washes, dosage, and brand-format dummies. Household i's valuation of these attributes is captured by taste parameters,

$$\beta_{it} = \beta + \beta_z z_{it} \tag{41}$$

To capture heterogeneous valuations of these attributes β_i has two components. The first is common to all households and is measured by a K-vector of parameters, β. The second captures the effect of household characteristics and is measured by an L-vector of parameters β_z.

Prices of product j in market t, p_{jt}, enter linearly into the indirect utility function. The marginal utility of income for each household is

$$\alpha_{it} = \alpha + \alpha_z z_{it} \tag{42}$$

In addition to a common component, α, the effect of different household characteristics on the marginal utility of income is measure by an L-vector of parameters, α_z.

To capture the impact of factors observed by the household, but not the econometrician, there is a household specific component of utility, ϵ_{ijt}. This error term is decomposed into two parts,

$$\epsilon_{ijt} = IC_{it} + \varepsilon_{ijt} \tag{43}$$

where IC_{it} are the unobserved inventory costs for household i in market t, and ε_{ijt} is an independently and identically distributed random utility shock that follows a member of the family of Generalised Extreme Value distribution.

Since the random shock is independent of all covariates, any endogeneity concerns arise from correlations between prices and inventory costs. As highlighted by Erdem et al. (2003), this is a prominent source of bias in storable good demand models and arises because inventories are unobserved by the researcher. Specifically, because both current prices and inventories are a function of past prices, the omission of inventories leads to price endogeneity. This problem could be resolved if instruments that are correlated with current prices but uncorrelated with past prices were available. However, since observed prices are serially correlated, finding such instruments is challenging.

Without instruments or a dynamic demand model, the source of bias cannot be corrected for fully.[28] Indeed, it is this source of bias I aim to reduce as much as possible by using accounting margins reported by firms.

4.2.2 Estimation and results

The demand model is estimated using maximum likelihood applied to Kantar purchase diary data described in Section 4.1.1.[29] It uses a sample of 100 purchases in each week from 1st January 2002 until 31st October 2012.[30]

Table 2 shows the results of the estimation of two choice models. The left column contains the parameter estimates of a logit specification and the right column shows the result of a nested logit model.

For the nested logit model, the set of SKUs sold in each week is partitioned into four groups based on the number of washes contained in each SKU: small (S), medium (M), large (L), and extra large (XL). The size boundaries of these groups correspond to the 25th, 50th, and 75th quantile of distribution of washes in each calendar year.

Both models include interactions between price and a measure of household income. Product characteristics include the size of the SKU purchased and the dosage - the amount of material (recommended) for use in a single wash. To control for household size,

[28] A dynamic demand model that integrates out over unobserved inventories. As noted above, this is likely to be infeasible within the constraints of an antitrust investigation.

[29] See Train (2009).

[30] The nested logit model is estimated sequentially. As highlighted by Train (2009) standard errors will be under-reported.

I also include the amount of washes purchased per equivalent adult in the household.[31] Detergent fixed effects are also included.

Table 2: Demand model parameter estimates

	Logit	**Nested Logit**
Price Params:		
Price	-0.437	-0.512
	(0.004)	(0.005)
Price x Income	0.217	0.337
	(0.009)	(0.012)
Characteristics:		
Washes	0.022	-0.007
	(0.001)	(0.003)
Washes per eq Ad.	-0.019	-0.010
	(0.002)	(0.006)
Dosage of Powder & Tabs	-2.334	-2.235
	(0.083)	(0.086)
Dosage of Liquids, Caps and Gel	-1.615	-1.480
	(0.075)	(0.078)
Other Params:		
Nesting Parameter		0.427
		(0.005)
Detergent Fixed Effects	Yes	Yes
N	56,200	56,200
Likelihood	-246,088	-244,576

As expected, price coefficients in both models are negative and households with higher income have a lower marginal utility of income. Though in the nested logit model, the price coefficient is more negative, but richer households are less price elastic.

Larger pack sizes are positively valued by households in the conditional logit model, especially in households with fewer people. However, in the nested logit model with size related choice sets, on average households tend to prefer smaller SKUs. However, the disutility from large SKU sizes per equivalent adult accrues more slowly.

The dosage, or amount of material needed to do a single wash is negatively valued, especially for 'solid' detergents. This is consistent with the fact that households value storage space. When the dosage is lower, households can store more washes without necessarily occupying more storage space. Indeed, as noted in Section 4.1, this is one of the driving factors behind the success of the new super-concentrated and gel detergent products. By itself, this suggest the presence of inter-temporal demand links through inventories - a source of mis-specification for this static demand model.

[31] To calculate equivalent adults, I use the OECD-modified scale. See http://www.oecd.org/eco/growth/OECD-Note-EquivalenceScales.pdf for details.

Table 3: Own-price elasticities

Brands	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012
Brand A	-5.91	-6.63	-6.71	-6.81	-7.05	-7.00	-7.14	-8.31	-8.55	-9.10	-8.82
Brand B	-5.32	-5.86	-5.55	-5.72	-5.75	-6.56	-6.99	-7.64	-7.80	-8.51	-8.19
Brand C	-4.63	-5.04	-5.06	-5.05	-5.60	-5.47	-5.64	-6.44	-5.93	-7.52	-8.28
Brand D	-5.65	-6.55	-6.29	-6.74	-7.23	-7.91	-7.87	-9.02	-9.20	-9.84	-9.13
Brand E	-6.32	-6.57	-6.50	-6.63	-6.63	-6.75	-7.21	-7.23	-7.55	-8.01	-8.34
Brand F	-4.07	-4.50	-4.31	-4.19	-3.93	-4.09	-5.12	-5.98	-5.99	-6.64	-6.65
PL	-4.18	-4.11	-3.77	-3.62	-3.77	-3.88	-4.51	-4.62	-4.59	-4.55	-5.04

Finally, the nesting parameter is 0.427 and is significantly different from 1. This indicates that that there are some unobserved correlations in the utility between detergents of similar sizes and rejects the independence of irrelevant alternatives imposed in the conditional logit. In subsequent analysis I use the nested logit specification.

The own-price elasticities of the 6 branded detergents produced by the two main firms and private label products (PL) are shown in Table 3. The static demand estimates of own-price elasticities typically lie between -5 and -10 for branded products

As I show in section 4.5, these price elasticities imply implausibly low margins of around 20 to 30 percent for a duopoly industry selling a highly differentiated products. This is suggestive of the aforementioned tendency of static demand models to produce overly elastic estimates of own-price elasticities in storable FMCG industries

The table also shows that the own-price elasticities are increasing over time. These increases in own-price elasticities are especially pronounced from 2008 onwards - a period that coincides with some of the most aggressive product compaction and intensifying price competition through discounting. This finding is discussed in more detail below.

4.3 Estimating inter-temporal profit ratios

In the context of an antitrust investigation brand level margins may be available over several years. However, in my case, I only have access to global, company-wide gross margins published in annual accounts. Therefore, I assume that the published gross margins in firm A's and firm B's annual accounts adjusted (if necessary) from 2002 to 2012 are a good approximation to the gross margins earned on sales of their laundry detergent portfolio in the UK.[32]

Plugging in the margins, using the Kantar data on industry outcomes in each year and combining with output of the demand model, the analyst searches for all pairs of

[32]The financial year for firm B starts midway through the year. As such, the annual report margins are adjusted to match calendar years in the data. Figures are omitted for confidentiality reasons.

parameters $(\psi_n^{own}, \psi_n^{cross})$ - the elements of $\bar{\Psi}_n$ - that satisfy

$$0 \leq 1 + \psi_n^{own} \leq 1 \tag{44}$$
$$1 \leq 1 + \psi_n^{cross} \leq 1 - \psi_n^{own} \tag{45}$$

and

$$\mu_n + \frac{\sum_{t \in y} q_{n,t}^\top \left(\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}_n\right)\right)^{-1} q_t}{\sum_{t \in y} q_{n,t}^\top p_t} = 0 \tag{46}$$

where $n \in \{f, y\}$, $f \in \{A, B\}$ and $y = 2002, \ldots, 2012$. [33]

Figure 6 shows the value of these parameters for both firms in each year in the data when the demand model outputs are evaluated at mean parameter values reported in Table 2.[34] It shows that the pairs of $(\psi_n^{own}, \psi_n^{cross})$ that satisfy equations (44-46) in each year for both firm's A and B lie on linear planes. The absolute value of the own and cross price elasticity biases are negatively correlated with one another and have similar slopes in each year. However, the intercept of the lines differ across year - especially for firm A.

The corresponding set of values that ψ^{own} can take for firm A and firm B in each year is shown in Figure 8. The blue band in the left panel of figure shows the range of values that ψ^{own} can take for firm A in each year from 2002 to 2012. The red band in the right panel shows this same information for firm B.

For firm A the figure suggests that the degree to which own-price demand derivatives are overstated by the static demand model increases over time. In 2002, the own-price elasticities are estimated to be overstated by 20% to 30%.

4.4 Product Compaction: Increase bias in elasticity estimates?

In line with the steady rate of product compaction of powder detergent, this bias in own-price elasticities increases gradually over time; by 2008 the estimated range is 25% to 30%. After 2008, when firm A introduced a new detergent, the degree to which price elasticities are overstated rapidly increases from 32% to 45% by 2010.

The sets of own-price elasticities once adjustments for estimates of the dynamic diversion ratios are shown in Figure 7. They show that own-price elasticities mostly lie between -3 and -6. This is much more elastic than the -5 to -10 range in Table 3. They also

[33]For the purpose of this section, I assume that the cartel in mainland Europe had no impact on the UK laundry detergent sector, and they priced detergent according to differentiated Bertrand competition between 2002 and 2005. I revisit this issue in Section 4.5.

[34]Construction of confidence sets requires calculating $(\psi^{own}, \psi^{cross})$ for the range of parameters values that lie in the corresponding confidence intervals in the demand model. With non-negligible additional computational costs, confidence sets can be implemented using a bootstrap procedure. At the time of writing, the confidence sets are yet to be calculated.

Figure 6: $(\psi^{own}, \psi^{cross})$ for Firm A and Firm B in each year from 2002 to 2012

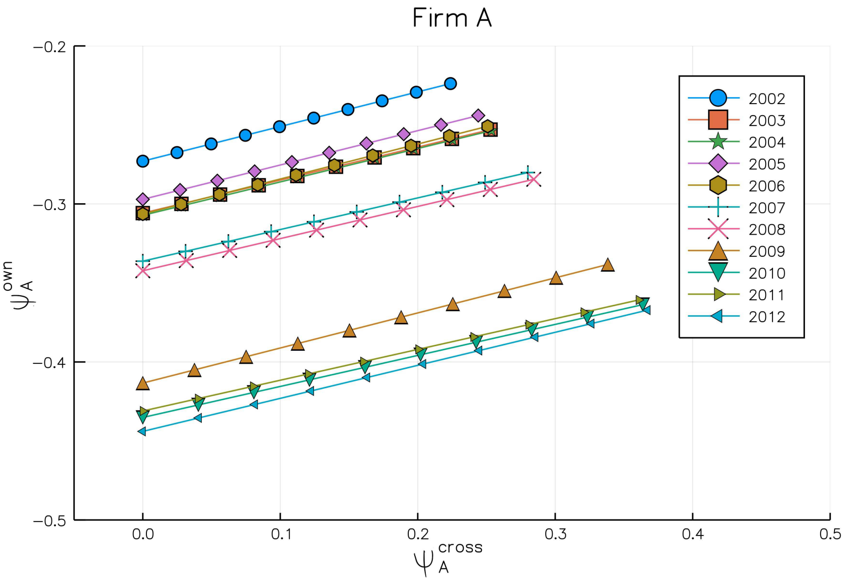

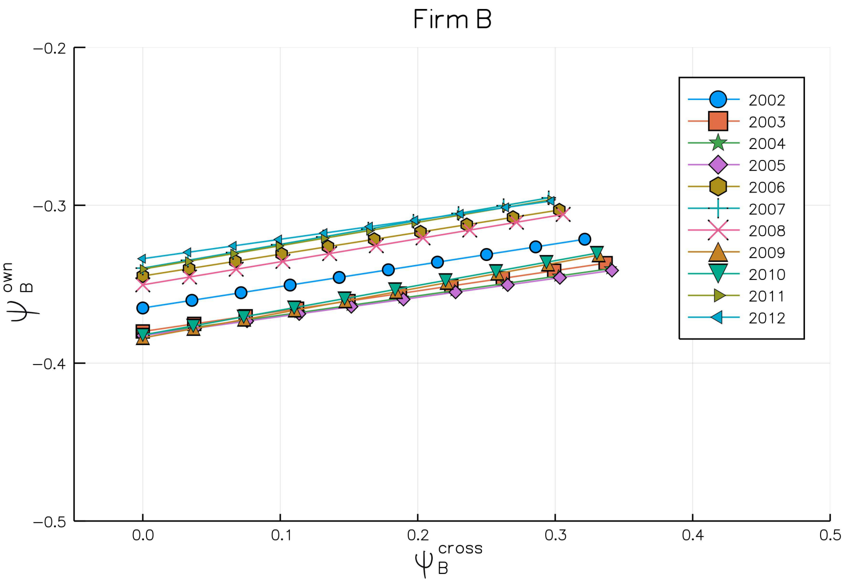

Figure 7: Own price elasticities by Brand: Unadjusted and Adjusted from 2002 to 2012

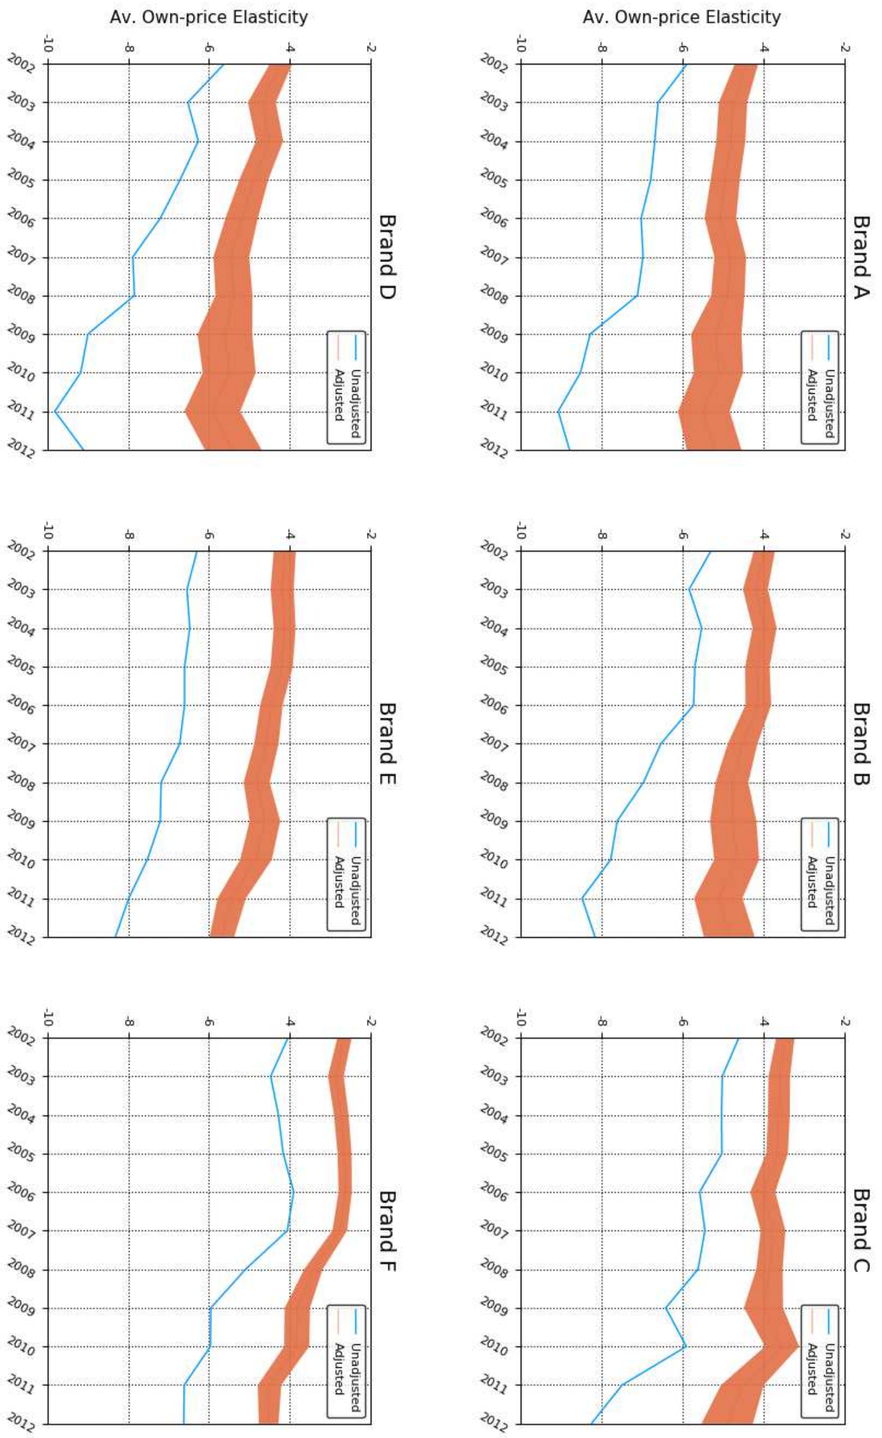

Figure 8: Estimated set for ψ^{own} for firm A and firm B, 2002 to 2012

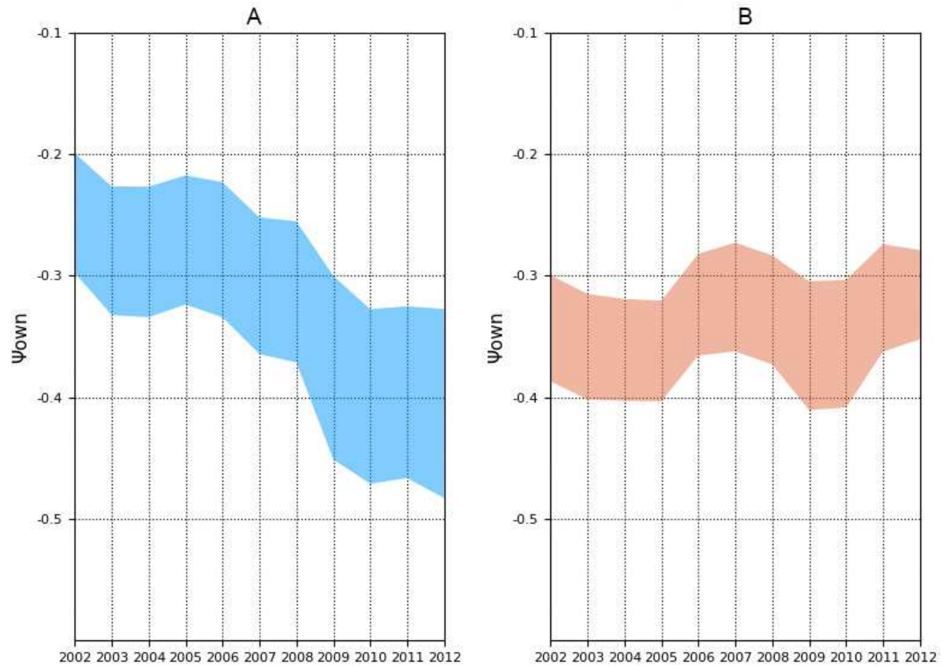

display much less pronounced declined over time as product are compacted and price discounting became more prevalent.

These above findings are consistent with the idea that demand dynamics changed as product innovation led to product compaction and storage costs for households fell. Consequently, as shown in Figure 4, households buy larger SKUs. In turn, households needed to buy less frequently and firms had fewer opportunities to cannibalise rivals' sales with their promotions. As a result discounts deepened and the intensity of competition increased (see Figure 5).

At first glance, the plot in the right panel of Figure 8 for firm B appears to cast some doubt onto this interpretation. For firm B, the sets of values taken by ψ^{own} is broadly constant over 2002 to 2012 and suggests that its own-price demand derivatives and elasticities are overstated by 30% and 40% by the static demand model. However, the estimates reflect changes in pricing strategies as well as the impact of product innovations.

Recall from Figure 5 that the pricing strategies of firm B and firm A diverge from around 2006. Namely, firm A increased the average posted price per wash so that the average price per wash paid by households remains broadly constant over time. In contrast, firm B held posted prices per wash approximately constant. As a result, deeper discounting for firm B led to lower purchase prices.

To isolate the impact product innovation has on the degree of bias in the own-price and cross-price demand derivatives, I divide through the range of values for ψ^{own} by the price index of purchased prices for each firm in each year. Figure 9 shows the set of values

Figure 9: Estimated set for ψ^{own} adjusted for observed purchased prices for firm A and firm B, 2002 to 2012

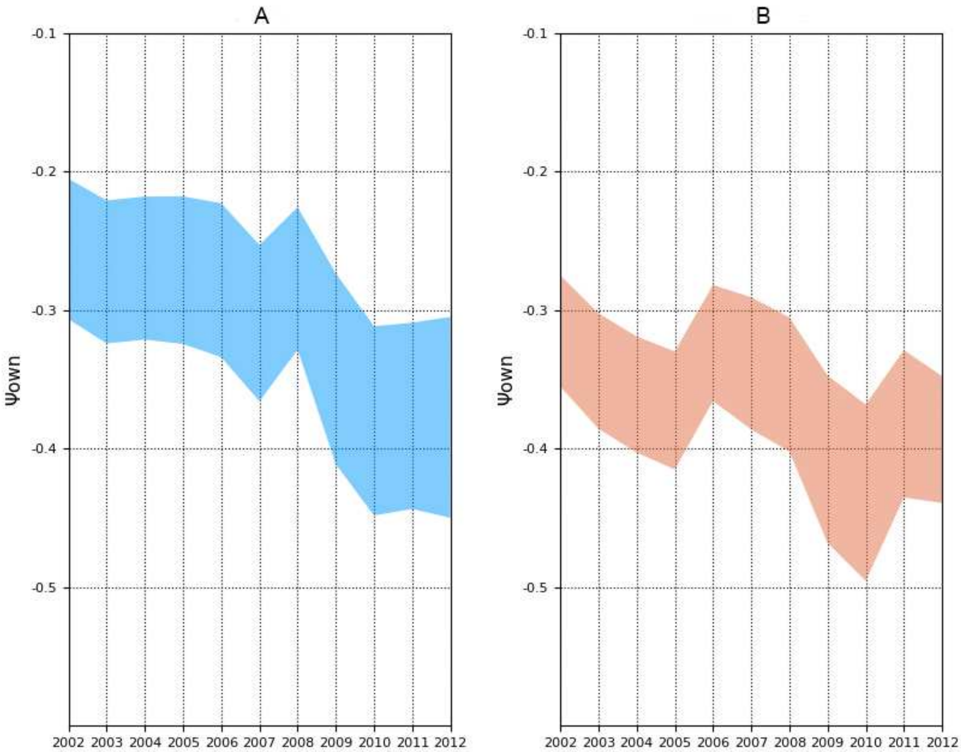

these adjusted estimates take on for firm A and firm B.

The figure shows that, because firm A purchase prices were relatively constant over time, the estimates of the bias are very similar to those in Figure 8. However, the estimate of the bias for firm B now more closely resembles the plot for firm A. One difference between the two figures is that the set of values for the bias increases from around 29% to 37% in 2006 to 35% to almost 50% in 2010. This coincides with the roll out of new product innovations.

Controlling for the changes in pricing behaviour, there is evidence from both firm A and firm B that innovation that lowered stocking costs exacerbates mis-specification of the static demand model. The result is increased bias in key quantities such as demand derivatives and elasticities that are important inputs into empirical antitrust policy.

4.5 The EU detergent cartel: UK impact?

In April 2011 the European Commission (EC) found that firm A, firm B and a firm C had entered into a cartel agreement that restricted competition in the market for heavy duty laundry detergent powder. The infringement was first brought to the attention of the EC when firm C 'blew the whistle' on the cartel in exchange for immunity from prosecution and/or reduced fines. Subsequent investigation led to the finding that the cartel was effective over the period 7th January 2002 to 8th March 2005. The EC highlighted four restrictive elements of cartel:

1. indirect price restrictions resulting from the parties agreeing not to pass on any cost savings that resulted from compaction of products
2. explicit reduction of promotional activity
3. direct price increases
4. an exchange of commercially sensitive information

According to the EC, the laundry detergent cartel (LD cartel) had anti-competitive effects in Belgium, France, Germany, Greece, Italy, Portugal, Spain and The Netherlands.[35] One notable absentee from the list of countries affected is the UK - a country in which the 'whistleblower', firm C, had virtually no market presence. However, the laundry detergent market in the UK is dominated by the other two firms in the cartel; firm A and firm B.

In this section I combine the structural demand model estimated in Section 4 together with various supply-side models to estimate market power under different types of firm conduct. By comparing the estimated margins associated with different models of firm behaviour to the observed margins for firm A and firm B from 2002 to 2005, I investigate the possibility that the collusive activities in mainland Europe were also evident in the UK.

I conduct two policy experiments. In the first I use the demand derivatives from the static demand model as the input in to an analysis of the market power of firm A and firm B during the cartel period. In the second, I repeat the analysis with the set of demand derivatives adjusted by the parameters $(\psi^{own}, \psi^{cross})$.

In the previous section, the set of estimates for $(\psi^{own}, \psi^{cross})$ for firm A and firm B in 2002 to 2005 assumed that the observed prices were the result of competitive behaviour. In this section, I wish to test whether this was the case. As such, the set of estimates for $(\psi^{own}, \psi^{cross})$ must be calibrated using another year where we know there was no cartel. Further, we also saw evidence that the product innovation from 2007 onwards had an impact on demand dynamics and, in turn, on the degree of bias of the estimates from the static demand model. Given these criteria, I use the set of parameters for $(\psi^{own}, \psi^{cross})$ from 2006.

In both experiments, demand model outputs are evaluated at mean parameter values. As a result the impact of uncertainty arising from estimation the demand model's parameters is not reflected in the predicted margins in the subsequent analysis. In a live policy setting producing confidence intervals that reflect uncertainty over the inputs into counterfactual simulations is necessary to ensure that the policy decisions are based on robust empirical evidence.[36]

[35] European Commission Decision, 'COMP/39579 - Consumer Detergents'.

[36] One way to implement this in the policy experiment described in this section is to bootstrap the demand model outputs over the range of parameter values in the confidence interval of interest. Using this approach, demand derivatives and purchase probabilities differ over bootstrap iterations and the first order conditions produce a corresponding set of predicted margins. However, in policy

However, since the goal of this section is to illustrate how the correcting elasticities for omitted demand dynamics can have a first order impact on policy analysis, I have not undertaken the additional computations needed to produce confidence sets around predicted margin intervals.

To estimate the market power, assumptions on the nature of supply-side competition and the shape of the cost function are added to each demand model in each experiment. Throughout I assume that marginal cost is (locally) constant.

To examine the impact of the intensity of competition on market power, the ownership matrix is altered to reflect different ownership structures. I consider four different supply side models to analyse market power: (1) firms engage in Bertrand-Nash price competition and each product is manufactured by a single firm, (2) assumes multi-product Bertrand-Nash competition, (3) assume firm A and firm B collude over powder products only,[37] and (4) assume that firm A and firm B set all prices jointly as a branded product monopolist. Scenarios (1) and (4) are intended as lower and upper bounds on the estimated market power.[38]

The interpretation of scenario (3) is also subject to the caveat that $\bar{\Psi}$ is held fixed under the assumption of multi-product Bertrand-Nash competition in 2006. Given the nature of the restrictive elements of the cartel listed above it is likely that the expected inter-temporal profit diversion ratio are likely to be reduced. As such, one might view the predicted margins under scenario (3) as a lower bound.

The results of the first experiment are shown in Figure 10. The left panel in the figure plots the observed margins in 2002 to 2005 for firm A alongside the margins implied by the demand model and each of the four firm conduct models. The right panel mirrors this analysis for firm B.

The figure shows that the observed margins of around 50% are well in excess of the margins implied by multi-product Bertrand-Nash. Moreover, the observed margins are well in excess of the 40% margin implied by monopoly pricing.

The fact that observed margins lie above the monopoly outcome serves as a warning that the demand model is mis-specified. Especially since the known biases that arise from omitting demand dynamics in storable goods industries would understate market power in this experiment. Based on these results from the mis-specified demand model, a policy analyst might conclude that the collusive conduct of these two dominant firms in Europe affected market outcomes in the UK laundry detergent market.

experiment 2, there is an additional step needed; within each bootstrap iteration the policy maker need to solve for $(\psi^{own}, \psi^{cross})$ in the chosen comparator market. Through this additional channel, uncertainty over demand model parameters can indirectly influence the size of the sets of predicted margins in counterfactual analysis.

[37]To reflect the cartel scenario, I assume firm B (firm A) take into account joint profits of a firm A (firm B) powder products, but do not set its price.

[38]I do not have information on the identity of producers of private label products. For the purposes of this paper I assume that firm A and firm B have no share in private label's profits. This is unlikely to be the case. As such the level of market power in multi-product scenarios (2) and (3) are likely to understate market power compared to the 'true' ownership matrix.

Figure 10: Policy experiment 1 ($\psi^{own} = \psi^{cross} = 0$): gross margins and implied market power during the cartel period: 2002 - 2005

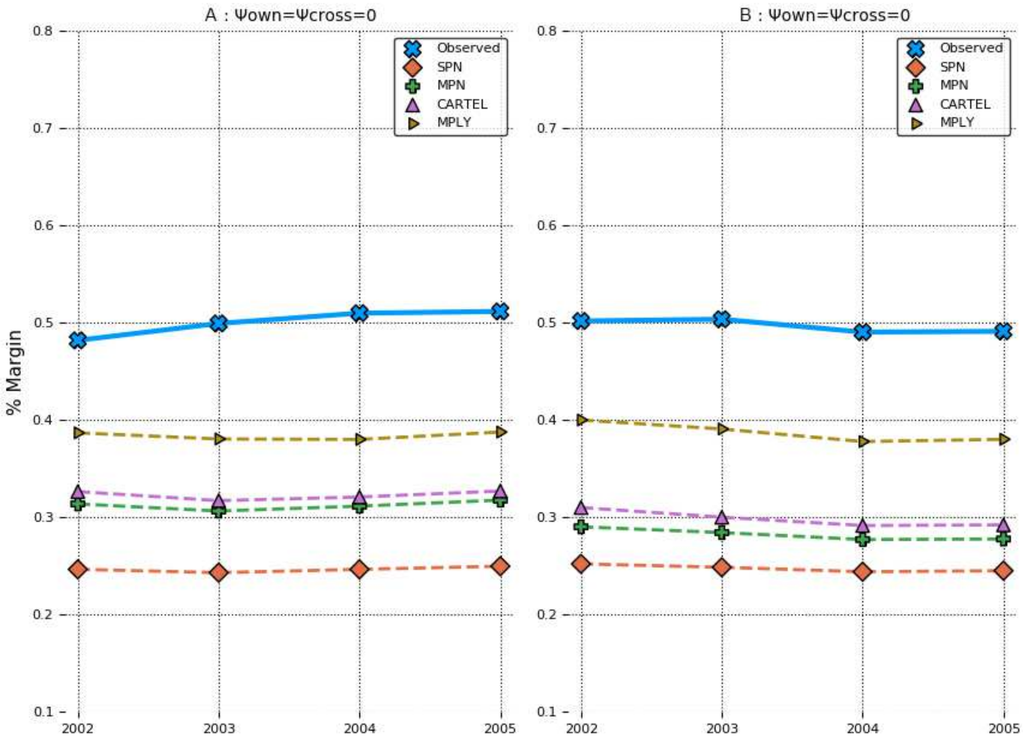

Note: SPN refers to single Product differentiated Nash-Bertrand competition, MPN refers multi-product differentiated Nash-Bertrand competition, CARTEL refers to pricing where firms A and B collude over prices for powdered detergents and MPLY refers to A and B pricing as a single entity (i.e. monopoly pricing).

Figure 11: Policy experiment 2: gross margins and implied market power during the cartel period: 2002 - 2005

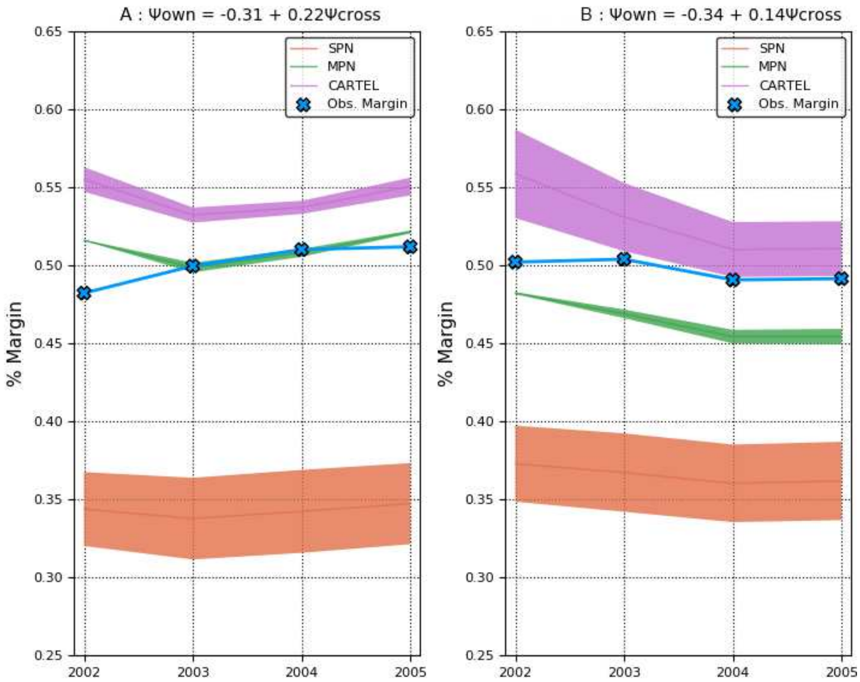

Note: SPN refers to single Product differentiated Nash-Bertrand competition, MPN refers multi-product differentiated Nash-Bertrand competition, CARTEL refers to pricing where firms A and B collude over prices for powdered detergents and MPLY refers to A and B pricing as a single entity (i.e. monopoly pricing).

Figure 11 shows the results of the same policy experiment but using the set of estimated parameters for ψ^{own} and ψ^{cross} from 2006 for both firms. The left panel of the figure plots the observed margin for firm A from 2002 to 2005 with the band of the predicted margins from the different models of conduct. Because the estimates of ψ^{own} and ψ^{cross} are set valued, the menu approach predicts a range of margins in each year. The right hand panel produces the same information for firm B. In both panels, the monopoly outcomes are omitted from the plot because they predict margins over 70%.

In contrast to the first policy experiment, the band of margins predicted under competitive conduct largely coincides with the observed margin for firm A in 2003 to 2005. In 2002, the observed margins lie below those implied by competitive conduct. For firm B, the predicted band of margins under competitive conduct lies below the observed margins. However, the observed margins also lie just below the lower bound of margins consistent with cartel conduct. While not conclusive, taken together the findings of this policy experiment do not find compelling evidence that the collusive conduct from mainland Europe occurred in the UK.

While it should be borne in mind that this experiment is based on an approximation to observed margins and does not use confidence sets that reflect demand parameter uncertainty, it shows how important it is to attempt to correct for known biases in empirical policy work.

5 Conclusion

In this paper I present an alternative approach to calculating price elasticities for antitrust policy that can be easily implemented within the timeframe of an antitrust investigation.

Combining the output of a static demand model estimated on weekly data with observed prices and firms' product-level accounting margins, I show how to recover parameters that capture the effect of missing demand dynamics on demand derivatives that are estimated by static models. I use these parameters to construct set-valued estimates of the price elasticity matrix that are suitable for use in policy simulations of consumer responses to permanent changes in firms' pricing behaviour.

The proposed approach makes use of the fact that margins are typically measured over a longer time horizon (i.e. year) than the period of analysis used to estimate the demand model (i.e. weekly market outcomes). As such they contain information on the aggregate impact of demand dynamics on market power over the reporting period. An advantage of this approach is that the severe bias associated with elasticities from a static model are reduced with minimal additional implementation costs.

This approach is applied to the UK laundry detergent industry. First, I explore the effect of product innovation on consumer demand dynamics, and the associated mis-specification bias of a static demand model.

Second, I conduct a policy experiment in which I assess whether anti-competitive conduct in mainland Europe's laundry detergent industry also occurred in the UK. I find that without using accounting margins to adjust elasticity estimates, policy simulations would suggest that observed margins are the product of anti-competitive conduct. However, when I use the set of bias-adjusted price elasticities, the analysis is much less clear-cut. If anything, there is little evidence that there were anti-competitive effects from the European laundry detergent cartel in the UK.

This paper demonstrates that the omission of demand dynamics has the potential to lead to misguided policy conclusions. The simple adjustment techniques presented here can help improve empirical information relied on by policy makers in lead contribute to improved policy decisions.

References

Bai, J. and Ng, S. (2008). Large dimensional factor analysis. *Foundations and Trends in Econometrics*, 3(2):89–163.

Bresnahan, T. (1987). Competition and collusion in the american automobile: The 1955 price war. *The Journal of Industrial Economics*, 35(4):pp457–482.

Crawford, A. (2017). High-dimensional inventories and consumer dynamics: demand estimation for fast moving consumer goods.

Erdem, T., Imai, S., and Keane, M. (2003). Brand and quantity choice dynamics under price uncertainty. *Quantitative Marketing and Economics*, 1:5–64.

Farrell, J. and Shapiro, C. (2010). Antitrust evaluation of horizontal mergers: An economic alternative to market definition. *The BE Journal of Theoretical Economics*, 10(1).

Frank, R. and Salkever, D. (1997). Generic entry and the pricing of pharmaceuticals. *Journal of Economics and Management Strategy*, Spring:75–90.

Hausman, J. A. and Leonard, G. K. (2002). The competitive effects of a new product introduction: A case study. *The Journal of Industrial Economics*, I(3):237–263.

Hendel, I. and Nevo, A. (2006). Measuring the implications of sales and consumer inventory behaviour. *Econometrica*, 74(6):1637–1673.

Jaffe, S. and Weyl, E. G. (2013). The first-order approach to merger analysis. *American Economic Journal: Microeconomics*, 5(4):188–218.

Miller, N. and Weinberg, M. (2017). Understanding the price effects of the miller/coors joint venture. *Econometrica*.

Moresi, S. (2010). The use of upward price pressure indices in merger analysis. *The Antitrust Source*, 9(3):8–9.

Mysliwski, M., Miessi Sanches, F., Silva Junior, D., and Srisuma, S. (2018). Implications of consumer loyalty for price dynamics when price adjustment is costly.

Nevo, A. (2001). Market power in the ready-to-eat cereal industry. *Econometrica*, 69:307–342.

Osborne, M. (2017). Approximating the cost-of-living index for a storable good. *Rotman School of Management Working Paper*.

Pires, T. (2016). Costly search and consideration sets in storable goods markets. *Quantitative Marketing and Economics*, 14(3):157–193.

Pittman, R. et al. (2009). Who are you calling irrational? marginal costs, variable costs, and the pricing practices of firms. Technical report, Department of Justice, Antitrust Division.

Rojas, C. (2008). Price competition in u.s. brewing. *The Journal of Industrial Economics*, 56(1):pp1–31.

Slade, M. (2004). Market power in uk brewing. *The Journal of Industrial Economics*, LII(1):133–163.

Sun, B. (2005). Promotion effect on endogenous consumption. *Marketing science*, 24(3):430–443.

Train, K. (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press, second edition.

Wang, E. Y. (2015). The impact of soda taxes on consumer welfare: implications of storability and taste heterogeneity. *The RAND Journal of Economics*, 46(2):409–441.

Annexes: Proofs and derivations

A.1 Proof of Proposition 1

The first order conditions for setting the price of product j is

$$\frac{\partial \pi_i}{\partial p_{j,t}} = q_{j,t} + \mathbb{E}_t \sum_{h=0}^{H} \sum_{k \in \mathcal{J}_f} \delta^h \left(\Delta_{t,t+h}\right)_{(j,k)} m_{k,t+h} = 0 \tag{47}$$

where $m_t := p_t - c_t$.

$$\begin{aligned}
\frac{\partial \pi_i}{\partial p_{j,t}} &= q_{j,t} + \mathbb{E}_t \sum_{h=0}^{H} \sum_{k \in \mathcal{J}_f} \delta^h \left[\Delta_{t,t+h}\right]_{(j,k)} \frac{\left[\Delta_{t,t}\right]_{(j,k)}}{\left[\Delta_{t,t}\right]_{(j,k)}} m_{k,t+h} \frac{m_{k,t}}{m_{k,t}} && (48)\\
&= q_{j,t} + \mathbb{E}_t \sum_{h=0}^{H} \sum_{k \in \mathcal{J}_f} \delta^h \left[\Delta^f_{t,t}\right]_{(j,k)} m_{k,t} \frac{\left[\Delta_{t,t+h}\right]_{(j,k)} m_{k,t+h}}{\left[\Delta_{t,t}\right]_{(j,k)} m_{k,t}} && (49)\\
&= q_{j,t} + \sum_{k \in \mathcal{J}_f} \left[\Delta_{t,t}\right]_{(j,k)} m_{k,t} \left(1 + \frac{\mathbb{E}_t \sum_{h=1}^{H} \delta^h \left[\Delta_{t,t+h}\right]_{(j,k)} m_{k,t+h}}{\left[\Delta_{t,t}\right]_{(j,k)} m_{k,t}} \right) && (50)\\
&= q_{j,t} + \sum_{k \in \mathcal{J}_f} \left[\Delta_{t,t}\right]_{(j,k)} m_{k,t} \left(1 + \left[\Psi_t\right]_{(j,k)}\right) && (51)
\end{aligned}$$

where

$$\left[\Psi_t\right]_{(j,k)} = \frac{\mathbb{E}_t \sum_{h=1}^{H} \delta^h \left[\Delta_{t,t+h}\right]_{(j,k)} m_{k,t+h}}{\left[\Delta_{t,t}\right]_{(j,k)} m_{k,t}} \; \forall \, j,k \in \mathcal{J}_f \tag{52}$$

is the dynamic diversion ratio.

Suppose A2 and A3 hold. Then there is a zero upper-bound on the diagonal elements of Ψ_t:

- If all products are gross substitutes the promoted good's future sales are a gross substitute for current sales. So a cut in the price today, will lead to a reduction in demand in the future. Since, all mark-ups are non-negative, the numerator and denominator in $\left[\Psi_t\right]_{(j,j)}$ have opposing signs, thus $\left[\Psi_t\right]_{(j,j)} \leq 0 \, \forall \, j \in \mathcal{J}_f$.

There is a zero lower-bound on the off-diagonal elements of Ψ_t:

- By a complementary logic the off-diagonal elements of Ψ_t are non-negative. This is because both current and future sales of rival product k are gross substitutes for the promoted good j. Under A2, both the numerator and denominator in $\left[\Psi_t\right]_{(j,k)}$ $j \neq k$ have the same sign, hence $\left[\Psi_t\right]_{(j,k)} \geq 0 \, \forall \, j \neq k, : j,k \in \mathcal{J}_f$.

Let

$$
\begin{aligned}
a_{j,k} &:= [\Delta_{t,t}]_{(j,k)} (p_{k,t} - c_{k,t}) \\
b_{j,k} &:= \mathbb{E}_t \sum_{h=1}^{H} \delta^h [\Delta_{t,t+h}]_{(j,k)} (p_{k,t+h} - c_{k,t+h})
\end{aligned}
$$

By A4,

$$
[\Delta_{t,t+h}]_{(j,k)} \leq 0 \text{ if } j = k, h = 0 \tag{53}
$$

$$
[\Delta_{t,t+h}]_{(j,k)} \geq 0 \text{ otherwise} \tag{54}
$$

for all $j, k \in \mathcal{J}_f$. Therefore

$$
a_{j,j} \leq 0, \; a_{j,k} \geq 0 \tag{55}
$$

$$
b_{j,k} \geq 0 \tag{56}
$$

for all $j, k \in \mathcal{J}_f$. Since

$$
[\Psi_t]_{(j,k)} := \frac{b_{j,k}}{a_{j,k}} \tag{57}
$$

Then for all product define the following bounds on inter-temporal profit ratios

$$
[\Psi_t]_{(j,j)} \leq 0 \tag{58}
$$

$$
[\Psi_t]_{(j,k)} \geq 0 \; j \neq k \tag{59}
$$

Further, from the first order conditions for product j

$$
q_{j,t} + \sum_{h=0}^{H} \sum_{k \in \mathcal{J}_f} \delta^h [\Delta_{t,t+h}]_{(j,k)} (p_{k,t+h} - c_{k,t+h}) = 0 \tag{60}
$$

$$
\implies a_{j,j} + b_{j,j} + \sum_{k \neq j} a_{j,k} + \sum_{k \neq j} b_{j,k} = -q_{jt} \leq 0 \tag{61}
$$

$$
\implies a_{j,j} + b_{j,j} \leq 0 \tag{62}
$$

$$
\implies b_{j,j} \leq -a_{j,j} \tag{63}
$$

$$
\implies \frac{b_{j,j}}{a_{j,j}} \geq -1 \tag{64}
$$

$$
\implies [\Psi_t]_{(j,j)} \geq -1 \tag{65}
$$

where in the third line above I use the fact $\sum_{k \neq j} a_{j,k} \geq 0$ and $\sum_{k \neq j} b_{j,k} \geq 0$ and in the second last line that $a_{j,j} \leq 0$. Thus, from eq (65) and (58) and adding 1, we have bounds for the diagonal elements of $1 + \Psi_t$:

$$
0 \leq 1 + [\Psi_t]_{(j,j)} \leq 1 \; \forall \, j \in \mathcal{J}_f \tag{66}
$$

Restriction A4 provides an upper bound on the off-diagonal elements of Ψ_t. Since $[\Psi_t]_{(j,j)} \leq 0$,

$$1 \leq 1 + [\Psi_t]_{(j,k)} \leq 1 - [\Psi_t]_{(j,j)} \;\forall\, j \neq k, j, k \in \mathcal{J}_f \tag{67}$$

Since the smallest value for $[\Psi_t]_{(j,j)} = -1$, then the maximum value for $[\Psi_t]_{(j,k)} = 1$. Therefore, a largest upper bound is given by,

$$1 \leq 1 + [\Psi_t]_{(j,k)} \leq 2 \,\forall\, j \neq k, j, k \in \mathcal{J}_f \tag{68}$$

■

A.2 Dynamic diversion ratio and price setting

Substituting the expression into equation (10) and rearranging

$$q_{j,t} + \sum_{k \in \mathcal{J}_f} [\Delta_{t,t}]_{(j,k)} m_{k,t} + \mathbb{E}_t \sum_{h=1}^{H} \sum_{k \in \mathcal{J}_f} \delta^h [\Delta_{t,t+h}]_{(j,k)} m_{k,t+h} = 0 \tag{69}$$

$$q_{j,t} + \sum_{k \in \mathcal{J}_f} [\Delta_{t,t}]_{(j,k)} m_{k,t} \left(1 + \frac{\mathbb{E}_t \sum_{h=1}^{H} \sum_{k \in \mathcal{J}_f} \delta^h [\Delta_{t,t+h}]_{(j,k)} m_{k,t+h}}{[\Delta_{t,t}]_{(j,k)} m_{k,t}} \right) = 0 \tag{70}$$

$$\implies q_{j,t} + \sum_{k \in \mathcal{J}_f} [\Delta_{t,t}]_{(j,k)} m_{k,t} \left(1 + [\Psi_t]_{(j,k)} \right) = 0 \tag{71}$$

A.3 Proof of Proposition 2

Below I derive the expression for the percentage margin of a group of n products over $t = 1, \ldots, \bar{T}$ periods.

$$
\begin{aligned}
q_t + \left(\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}\right)\right)(p_t - c_t) &= 0 && (72)\\
\implies p_t - c_t &= -\left(\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}\right)\right)^{-1} q_t && (73)\\
\implies q_{n,t}^{\top}(p_t - c_t) &= -q_{n,t}^{\top}\left(\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}\right)\right)^{-1} q_t && (74)\\
\implies \bar{T}^{-1} \sum_{t=1}^{\bar{T}} q_{n,t}^{\top}(p_t - c_t) &= -\bar{T}^{-1} \sum_{t=1}^{\bar{T}} q_{n,t}^{\top}\left(\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}\right)\right)^{-1} q_t && (75)\\
\implies \frac{\bar{T}^{-1} \sum_{t=1}^{\bar{T}} q_{n,t}^{\top}(p_t - c_t)}{\bar{T}^{-1} \sum_{t=1}^{\bar{T}} q_{n,t}^{\top} p_t} &= -\frac{\bar{T}^{-1} \sum_{t=1}^{\bar{T}} q_{n,t}^{\top}\left(\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}\right)\right)^{-1} q_t}{\bar{T}^{-1} \sum_{t=1}^{\bar{T}} q_{n,t}^{\top} p_t} && (76)\\
\implies \mu_{n,[1:\bar{T}]} &= -\frac{\sum_{t=1}^{\bar{T}} q_{n,t}^{\top}\left(\Delta_{[t,t]} \circ \left(1 + \bar{\Psi}\right)\right)^{-1} q_t}{\sum_{t=1}^{\bar{T}} q_{n,t}^{\top} p_t} && (77)
\end{aligned}
$$

where $\mu_{n,[1:\bar{T}]}$ is the observed gross margin of products in group n over $t = 1, \ldots, \bar{T}$ periods.

A.4 Non-singularity of $\Delta_{[t,t]} \circ \left(1+\bar{\Psi}\right)$

Define weights $\tilde{\omega}$ where $\tilde{\omega}_j = 1$, so that product j's weighted long-run aggregate diversion ratio equals 1.

$$[\Delta_{t,t}]\left[1+\bar{\Psi}\right]_{(j,j)} = -\sum_{k\neq j}\tilde{\omega}_k\left[\Delta_{t,t}\right]_{(j,k)}\left(1+\left[\bar{\Psi}\right]_{(j,k)}\right) \tag{78}$$

$$\implies 1 = \frac{-\sum_{k\neq j}\tilde{\omega}_k\left[\Delta_{t,t}\right]_{(j,k)}\left(1+\left[\bar{\Psi}\right]_{(j,k)}\right)}{\left[\Delta_{t,t}\right]_{(j,j)}\left(1+\left[\bar{\Psi}\right]_{(j,j)}\right)} \tag{79}$$

$$= -\sum_{k\neq j}\tilde{\omega}_k\frac{\left[\Delta_{t,t}\right]_{(j,k)}\left(1+\left[\bar{\Psi}\right]_{(j,k)}\right)}{\left[\Delta_{t,t}\right]_{(j,j)}\left(1+\left[\bar{\Psi}\right]_{(j,j)}\right)} \tag{80}$$

Then for all other products, $j' \neq j$, the matrix $\Delta_{[t,t]} \circ \left(1+\bar{\Psi}\right)$ is non-singular if

$$1 \neq -\sum_{k\neq j'}\frac{\tilde{\omega}_k}{\tilde{\omega}_{j'}}\frac{\left[\Delta_{t,t}\right]_{(j',k)}\left(1+\left[\bar{\Psi}\right]_{(j,k)}\right)}{\left[\Delta_{t,t}\right]_{(j',j')}\left(1+\left[\bar{\Psi}\right]_{(j,j)}\right)}$$

That is, no other products' weighted long-run aggregate diversion ratio equals 1. As such, non-singularity of the matrix is linked to the uniqueness of long-run substitution of products produced by that firm.

With assumptions A5 and A6, this condition can be express in terms of weighted short-run diversion ratios.

$$\frac{(1+\psi_n^{own})}{(1+\psi_n^{cross})} = \textstyle\sum_{k\neq j'}\tilde{\omega}_k DR_{(j,k)} \neq \displaystyle\sum_{k\neq j'}\frac{\tilde{\omega}_k}{\tilde{\omega}_{j'}}DR_{(j',k)}$$

A.5 Price forecasting model

When choosing the current price, the firms takes into account the effect of a change in price today on expected prices in the future. Firms believe that the price evolves according to the following stationary factor model that is known to all firms prior to setting prices

$$\begin{aligned} p_t &= \lambda_j f_t + \epsilon_{jt} \\ f_{t+1} &= A f_t + u_{t+1} \end{aligned}$$

where

$$\left[\begin{array}{c}\epsilon_t \\ u_{t+1}\end{array}\right] \sim N\left(\left[\begin{array}{c}0\\0\end{array}\right], \left[\begin{array}{cc}\Sigma_{pp} & 0 \\ 0 & \Sigma_{ff}\end{array}\right]\right)$$

A.5.1 One-period ahead

Stacking products and rolling forward one period

$$
\begin{aligned}
p_t &= \Lambda f_t + \epsilon_t \\
f_{t+1} &= A f_t + u_{t+1}
\end{aligned}
$$

where $\Lambda = [\lambda_1, \ldots, \lambda_J]$ is a $J \times R$ matrix of time-invariant factor loadings and R is the number of factors.

The covariance matrix is

$$
\mathbb{C}_t \left(\begin{bmatrix} p_t \\ f_{t+1} \end{bmatrix} | f_t \right) = \begin{bmatrix} \Omega_{pp} & \Omega_{pf} \\ \Omega_{fp} & \Omega_{ff} \end{bmatrix}
$$

For brevity suppress conditioning variables, then the components of the covariance matrix are

$$
\begin{aligned}
\underset{J\times J}{\Omega_{pp}} &= \mathbb{E}_t \left[(\Lambda f_t + \epsilon_t) (\Lambda f_t + \epsilon_t)^\top \right] \\
&= \Lambda f_t f_t^\top \Lambda^\top + \Sigma_{pp}
\end{aligned}
$$

$$
\begin{aligned}
\underset{R\times R}{\Omega_{ff}} &= \mathbb{E}_t \left[(A f_t + u_{t+1}) (A f_t + u_{t+1})^\top \right] \\
&= A f_t f_t^\top A^\top + \Sigma_{ff}
\end{aligned}
$$

$$
\begin{aligned}
\underset{J\times R}{\Omega_{pf}} &= \mathbb{E}_t \left[(\Lambda f_t + \epsilon_t) (A f_t + u_{t+1})^\top \right] \\
&= \Lambda f_t f_t^\top A^\top
\end{aligned}
$$

$$
\begin{aligned}
\underset{R\times J}{\Omega_{fp}} &= \mathbb{E}_t \left[(A f_t + u_{t+1}) (\Lambda f_t + \epsilon_t) \right] \\
&= A f_t f_t^\top \Lambda^\top
\end{aligned}
$$

So

$$
f_{t+1} | p_t = \tilde{p}, f_t \sim N(\bar{\mu}, \bar{\Sigma})
$$

where

$$
\begin{aligned}
&\bar{\mu}_1 = A f_t + \Omega_{fp} \Omega_{pp}^{-1} (\tilde{p} - \Lambda f_t) \\
&\bar{\Sigma}_1 = \Omega_{ff} - \Omega_{fp} \Omega_{pp}^{-1} \Omega_{pf}
\end{aligned}
$$

So

$$
\begin{aligned}
\mathbb{E}_t\left[p_{t+1}|p_t\right] &= \Lambda\mathbb{E}_t\left[f_{t+1}|p_t\right] + \mathbb{E}_t\left[\epsilon_{t+1}|p_t\right] \\
&= \Lambda A f_t + \Lambda\Omega_{fp}\Omega_{pp}^{-1}\left(p_t - \Lambda f_t\right)
\end{aligned}
$$

Then

$$
\begin{aligned}
\mathbb{E}_t\left[\frac{dp_{t+1}}{dp_t}\right] &= \Lambda\Omega_{fp}\Omega_{pp}^{-1} \\
&= \Lambda A f_t f_t^\top \Lambda^\top \left[\Lambda f_t f_t^\top \Lambda^\top + \Sigma_{pp}\right]^{-1}
\end{aligned}
$$

which is the expectation of a matrix of constants.

A.5.2 h-periods ahead

Stacking products and rolling forward h periods

$$
\begin{aligned}
p_t &= \Lambda f_t + \epsilon_t \\
f_{t+h} &= A^h f_t + \sum_{r=0}^{h-1} A^r u_{t+r+1}
\end{aligned}
$$

where $\Lambda = [\lambda_1, \ldots, \lambda_J]$ is a $J \times R$ matrix of time-invariant factor loadings and R is the number of factors.

The covariance matrix

$$
\mathbb{C}_t\left(\left[\begin{array}{c} p_t \\ f_{t+h} \end{array}\right]|f_t\right) = \left[\begin{array}{cc} \Omega_{pp} & \Omega_{pf}^h \\ \Omega_{fp}^h & \Omega_{ff}^h \end{array}\right]
$$

For brevity suppress conditioning variables, then the components of the covariance matrix are

$$
\begin{aligned}
\underset{J\times J}{\Omega_{pp}} &= \mathbb{E}_t\left[\left(\Lambda f_t + \epsilon_t\right)\left(\Lambda f_t + \epsilon_t\right)^\top\right] \\
&= \Lambda f_t f_t^\top \Lambda^\top + \Sigma_{pp}
\end{aligned}
$$

$$
\begin{aligned}
\underset{R\times R}{\Omega_{ff}^h} &= \mathbb{E}_t\left[\left(A^h f_t + \sum_{r=0}^{h} A^r u_{t+r+1}\right)\left(A^h f_t + \sum_{r=0}^{h-1} A^r u_{t+r+1}\right)^\top\right] \\
&= A^h f_t f_t^\top A^{h,\top} + \sum_{r=0}^{h-1} A^r \Sigma_{ff} A^{r,\top}
\end{aligned}
$$

$$
\begin{aligned}
\underset{J\times R}{\Omega^h_{pf}} &= \mathbb{E}_t\left[(\Lambda f_t+\epsilon_t)\left(A^h f_t+\sum_{r=0}^{h-1}A^r u_{t+r+1}\right)^\top\right]\\
&= \Lambda f_t f_t^\top A^{h,\top}
\end{aligned}
$$

$$
\begin{aligned}
\underset{R\times J}{\Omega^h_{fp}} &= \mathbb{E}_t\left[\left(A^h f_t+\sum_{r=0}^{h-1}A^r u_{t+r+1}\right)(\Lambda f_t+\epsilon_t)\right]\\
&= A^h f_t f_t^\top \Lambda^\top\\
&= A^{h-1}\Omega^1_{fp}
\end{aligned}
$$

So

$$
f_{t+h}|p_t=\tilde{p}, f_t \sim N(\bar{\mu}^h, \bar{\Sigma}^h)
$$

where

$$
\begin{aligned}
\bar{\mu}^h &= A^h f_t + \Omega^h_{fp}\left(\Omega_{pp}\right)^{-1}\left(\tilde{p}-\Lambda f_t\right)\\
\bar{\Sigma}^h &= \Omega^h_{ff} - \Omega^h_{fp}\left(\Omega_{pp}\right)^{-1}\Omega^h_{pf}
\end{aligned}
$$

So

$$
\begin{aligned}
\mathbb{E}_t\left[p_{t+h}|p_t\right] &= \Lambda\mathbb{E}_t\left[f_{t+h}|p_t\right]+\mathbb{E}_t\left[\epsilon_{t+h}|p_t\right]\\
&= \Lambda A^h f_t + \Lambda\Omega_{fp}\left(\Omega^h_{pp}\right)^{-1}\left(p_t-\Lambda f_t\right)
\end{aligned}
$$

Then

$$
\begin{aligned}
\mathbb{E}_t\left[\frac{dp_{t+h}}{dp_t}\right] &= \Lambda\Omega^h_{fp}\left(\Omega_{pp}\right)^{-1} = \Lambda A^{h-1}\Omega^1_{fp}\left(\Omega_{pp}\right)^{-1}\\
&= \Lambda A^h f_t f_t^\top\Lambda^\top\left[\Lambda f_t f_t^\top \Lambda^\top + \Sigma_{pp}\right]^{-1}
\end{aligned}
$$

which is the expectation of a matrix of constants