Product Market Dynamics over the Business Cycle

Geumbi Park*
Texas A&M University

November 19, 2025

[Link to the latest version]

Abstract

Business cycles reshape product market dynamics, but these responses differ markedly across markets. This heterogeneity arises from differences in consumer demand structures and their interaction with producers' entry and exit decisions—an aspect largely overlooked in the literature. This paper examines how product entry and exit over the business cycle depend on the structure of consumer demand. I develop a model of strategic firm dynamics under imperfect competition with heterogeneous producers and consumers, allowing demand composition to shape the evolution of market structure. Using detailed scanner data from the consumer goods sector, I estimate the model separately for each differentiated product market to analyze producer responses during the Great Recession. I find that negative aggregate demand shocks lead firms to reduce entry and increase exit, with the magnitude of these responses varying systematically with the price elasticity of demand. Markets with less elastic demand experience larger contractions in product variety, and interactions between elasticity and product characteristics further amplify recessionary effects. I also show that the welfare costs of the business cycle are substantially larger when accounting for demand heterogeneity, with most losses stemming from declines in consumer surplus. In the absence of consumer heterogeneity, the welfare loss is significantly understated. Overall, the results demonstrate that consumer heterogeneity is a key driver of how recessions affect market structure and welfare.

*Email: geumbi.park@tamu.edu. I am indebted to my advisors, Sarah Zubairy, Pedro Bento, and Fernando Luco, for their invaluable advice and support throughout this project. I thank Hie Joo Ahn, Zach Bethune, Yongsung Chang, Colin Hottman, Aljoscha Janssen, Yoon Joo Jo, Bariş Kaymak, Aubhik Khan, Ryan Kim, Ben Klopack, Annie Lee, Maximiliano Machado, Robert Minton, Sara Moreira, Jake Orchard, Ekaterina Peneva, Erick Sager, Francisca Sara-Zaror, Tatevik Sekhposyan, Rohan Shah, Simon Smith, Dongho Song, Martin Souchier, Brad Strum, Daniel Villar, Choongryul Yang, Yangkeun Yun as well as seminar and conference participants at Texas A&M, SEA Annual Meeting, ASSA Annual Meeting, North America Summer Meeting of the Econometric Society, Federal Reserve Board, Boston College PhD Conference, Texas Macro Job Candidate Conference, Fall Midwest Macroeconomics Meeting, for discussions and comments. I extend my special thanks to economists at the Federal Reserve Board for their insightful feedback. Portions of this research were conducted with the advanced computing resources provided by Texas A&M High Performance Research Computing. This study contains my own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1 Introduction

How does demand heterogeneity shape the impact of the business cycle through product entry and exit? At the aggregate level, fluctuations in entry and exit generate cyclicality in product variety (Broda and Weinstein, 2010; Argente, Lee, and Moreira, 2018), but this relationship masks substantial heterogeneity in cyclicality across products. For example, following the Great Recession, the number of products in the canned soup market fell by about 15%, whereas markets for jams and bottled water experienced only modest declines of around 1.5%. These differences suggest that recessions do not uniformly affect the product space, leading to systematic variation in welfare impacts across products. Yet relatively little attention has been paid to understanding these micro-level differences.

To fill this gap, this paper develops a unified structural and empirical framework that links heterogeneous consumer demand to producers' dynamic entry and exit decisions over the business cycle. In my framework, demand heterogeneity arises because consumers differ in their sensitivity to prices and in their tastes for different products, while products themselves vary in their characteristics. This is essential for explaining why recessions generate heterogeneous impacts across the product space and for accurately assessing the welfare costs of the business cycle.

My approach proceeds as follows. First, I build a structural model that embeds demand heterogeneity into producers' entry, exit, and market coverage decisions over the business cycle, where coverage is measured by the number of stores in which each product is sold. This framework links demand-side variation to firms' dynamic choices, allowing me to study how differences in consumer preferences translate into heterogeneous product-market dynamics. Second, I use a granular product-level dataset covering hundreds of differentiated goods, capturing rich demand heterogeneity and detailed patterns of entry and exit across industries.[1] Third, I estimate the model separately for each industry to recover the parameters governing consumer demand and producers' dynamic decisions, which reveal the underlying economic differences across markets. This estimation also recovers a key state variable, *appeal*, which summarizes the demand shifter for each product.[2] Finally, I quantify how the interaction between demand heterogeneity and producers'

[1]Throughout the paper, I use the terms *industry* and *market* interchangeably. An industry is defined as the market for a specific good (e.g., cereal, yogurt, bottled water), which consists of a set of products.

[2]Following Hottman, Redding, and Weinstein (2016), I use the term "appeal" rather than "quality" to remain agnostic about whether demand variation reflects vertical quality, taste, or other unobserved factors beyond price and observable

dynamic choices amplifies cyclical welfare losses—costs that are substantially understated when demand heterogeneity is ignored.

Analyzing the interaction while accounting for heterogeneity poses two main challenges. First, studying product-level dynamics requires rich microdata that reveal heterogeneity obscured at the aggregate level. However, the granularity of such data also makes it difficult to integrate into a macroeconomic framework. Second, entry and exit decisions are dynamic and strategic: they must account for competitors' states and actions when making their own decisions. This interdependence creates a high-dimensional state space and substantial computational complexity.

In this paper, I address these challenges by integrating tools from both macroeconomics and industrial organization. Specifically, I incorporate a canonical discrete-choice demand system from industrial organization (Berry, Levinsohn, and Pakes, 1995) into a model of strategic firm dynamics, introducing an aggregate demand shock to capture business cycle fluctuations. This integrated framework offers two key advantages. First, it allows for flexible substitution grounded in microdata through consumer heterogeneity, enabling novel micro-level empirical analysis that cannot be achieved in models assuming a representative demand structure (Nevo, 2011). Second, the framework assumes that heterogeneous producers do not fully observe competitors' detailed states or decisions. Rather, they base their decisions on an aggregate moment summarizing overall profitability and competition (Ifrach and Weintraub, 2017). This structure makes dynamic strategic interactions tractable while preserving the key strategic forces, simplifying computation.[3]

Using this framework, I provide three key findings. First, a one-standard-deviation decrease in the median price elasticity of demand (in absolute value) is associated with a 14-percentage-point larger decline in product entry across industries during the Great Recession. This is consistent with higher estimated sunk entry costs in markets with less elastic demand, implying greater difficulty of entry during recessions. Within industries, exit and market coverage adjustments also vary systematically, with effects concentrated among low-appeal products that are more likely to exit and experience larger reductions in market coverage. Evidence from both the extensive and intensive

characteristics. This concept of product appeal also differs from productivity, which refers to a firm's ability to produce more output given the same inputs.

[3] I adopt the Moment-Based Markov Equilibrium framework of Ifrach and Weintraub (2017), which provides a tractable approximation to the Markov Perfect Equilibrium in dynamic oligopoly (e.g., Ericson and Pakes, 1995). The idea follows Krusell and Smith (1998) and Hopenhayn (1992), where aggregate moments summarize relevant state variables to simplify agents' dynamic optimization.

margins shows that these effects are stronger for products facing less elastic demand, except for those with sufficiently high appeal to remain relatively resilient. Low-appeal, niche products are disproportionately vulnerable to exit, and this is further amplified with less elastic demand.[4]

Second, the welfare costs of business cycles driven by aggregate demand shocks are substantial. The estimated welfare loss is about 1.5% of pre-recession personal consumption expenditure (PCE) on food products (approximately $37 billion). This magnitude contrasts with the minuscule welfare loss in the conventional benchmark of Lucas (1987), as my analysis focuses on micro-level product-market entry and exit dynamics and demand heterogeneity. These losses are concentrated in consumer surplus because the remaining high-appeal products in recessions charge higher prices, limiting substitution toward cheaper, low-appeal varieties. This highlights that the welfare gains from access to high-appeal products are outweighed by the losses consumers incur from higher prices. The decline is also significantly larger in markets with less elastic demand, where substitution is more limited. Given that consumer surplus accounts for roughly 10% of total social welfare, these findings indicate that consumers are disproportionately affected by recessions through changes in product-market dynamics.

Lastly, incorporating demand heterogeneity increases the quantified welfare costs of recessions by about 5 percent. This result comes from a counterfactual analysis comparing welfare outcomes in an economy with demand heterogeneity to one without. The latter features a multinomial logit demand system, which closely approximates a constant elasticity of substitution (CES) demand structure (Anderson, De Palma, and Thisse, 1992), thereby limiting flexible substitution across heterogeneous products. This finding highlights that conventional frameworks mute welfare impacts from product reallocation and substantially understate consumer losses, by about 0.2 percent of average PCE on food products.

This paper builds on highly disaggregated NielsenIQ Retail Scanner data, where my analysis focuses on about 15 consumer goods industries over the period 2006–2019. The sample period spans the Great Recession, providing a natural setting to study how large aggregate demand shocks influence product market dynamics. The granularity of the dataset allows me to link micro-level demand heterogeneity to dynamic producer behavior through observations of product-level sales, entry, exit, and market coverage adjustments.

[4]The inverse relationship between elasticity and market power (e.g., Lerner index) can be misleading in this context.

The structural model in this paper captures the behavior of both consumers and producers. On the consumer side, they are heterogeneous in their sensitivity to prices, which I model as a random coefficient drawn from a distribution characterized by mean and variance parameters. The model also accounts for unobserved demand shifters in consumer utility. Consumers value product availability, which corresponds to producers' predetermined market coverage decisions. On the producer side, the model features heterogeneous and strategic firms that make dynamic decisions regarding entry, exit, and market coverage adjustments over time. Producers operate under aggregate demand shocks that affect the overall market size each firm faces and interact strategically through aggregate moments rather than tracking the full distribution of competitors' states.

My analysis includes the estimation of the structural model, conducted in two steps. First, I estimate a random-coefficients logit demand system (Berry, Levinsohn, and Pakes, 1995). Because estimating this model is computationally intensive, I adopt the approximation approach developed by Salanié and Wolak (2022). I apply this procedure separately to each industry and use instrumental variables introduced by Gandhi and Houde (2019) to identify the demand parameters, from which I recover the price elasticity of demand. Second, I estimate the cost parameters governing product entry, exit, and market coverage adjustments over the business cycle. I then estimate the model with a full-solution maximum-likelihood approach (Rust, 1987; Igami, 2017). The estimation compares model-implied policies with observed behavior in the microdata.

Finally, I conduct a welfare analysis and perform counterfactual exercises using the estimated parameters. I first compare welfare across different scenarios, with and without aggregate demand shocks. This exercise relies on the same estimated demand and dynamic parameters for each industry. These parameters are recovered from observed micro-level behavior in the data, rather than calibrated to match empirical moments, ensuring that the welfare comparisons directly reflect the observed structure of the economy. I then conduct an additional counterfactual that contrasts the baseline economy with a case without demand heterogeneity. I set the variance of the random coefficients in the demand system to zero, so that all consumers within each industry share the same price sensitivity. The model still allows for market-level differences in price sensitivity, resembling a nested CES demand structure commonly used in macroeconomic analysis.

In sum, this paper highlights that consumer demand heterogeneity is central to understanding how recessions reshape product markets. By linking micro-level product dynamics to aggregate business cycle fluctuations, the analysis shows that demand structure plays a key role in explaining changes in entry and exit, and that welfare losses are amplified through changes in market structure. It underscores the importance of incorporating demand heterogeneity into policy evaluations of business cycle costs, as ignoring it substantially understates welfare losses.

Related Literature. This paper contributes to three strands of the literature. First, it bridges two areas in macroeconomics: the literature on entry and exit over the business cycle and the literature emphasizing the role of demand structure in shaping aggregate outcomes.

A large body of work on business cycle dynamics highlights the role of firm, establishment, and product entry and exit in aggregate fluctuations (Siemer, 2014; Lee and Mukoyama, 2015; Clementi and Palazzo, 2016; Hamano and Zanetti, 2017; Gamber, 2023).[5] Entry and exit influence aggregate fluctuations through changes in market structure (Jaimovich and Floetotto, 2008; Etro and Colciago, 2010) and through love-of-variety preferences (Bilbiie, Ghironi, and Melitz, 2012). However, while these models often incorporate firm heterogeneity, they typically abstract from consumer heterogeneity, leaving the interaction between demand-side variation and firms' entry and exit decisions unexplored. My paper jointly accounts for these two dimensions by examining product-level heterogeneity observed in the microdata.

In the literature examining how demand structures shape macroeconomic outcomes over the business cycle, the main channels operate through cyclical changes in consumers' behavior (Jaimovich, Rebelo, and Wong, 2019; Argente and Lee, 2021; Michelacci, Paciello, and Pozzi, 2022) or through differences in preferences and consumption patterns between necessities and other goods (Bils and Klenow, 1998; Charles and Stephens Jr, 2006; Orchard, 2022; Becker, 2024). My findings connect to this literature by showing that the impact of the business cycle depends on the price elasticity of demand: industries producing less price-elastic (often necessity) goods experience larger cyclical effects. I extend this literature by jointly modeling product-level demand heterogeneity and producers' dynamic entry and exit decisions.

Second, this paper relates to the classic macroeconomic literature on the welfare costs of business cycles. This literature originates from the conventional benchmark of Lucas (1987) and

[5] In cases involving multi-product firms, adjustments in entry and exit often take the form of changes in product scope, a margin that I abstract from in this paper.

has been extended to incorporate heterogeneous agents and incomplete markets (Imrohoroğlu, 1989; Atkeson and Phelan, 1994; Krusell and Smith Jr, 1999; Storesletten, Telmer, and Yaron, 2001; Krebs, 2003, 2007; De Santis, 2007), as well as alternative preference specifications (Obstfeld, 1994; Dolmas, 1998; Otrok, 2001). It has also advanced in measuring and quantifying welfare costs (Epaulard and Pommeret, 2003; Alvarez and Jermann, 2004; Chatterjee and Corbae, 2007; Gali, Gertler, and Lopez-Salido, 2007; Reis, 2009; Constantinides, 2025). My paper contributes to this literature by developing a framework that accounts for heterogeneous consumers and producers, incorporates product market dynamics, and estimates welfare using detailed micro-level data.

Lastly, this paper connects the industrial organization literature with macroeconomics. Across various industry contexts, it has been well documented that the introduction of new products, as well as firms' entry and exit, affects competition and welfare (Nevo, 2001; Seim, 2006; Dunne et al., 2013; Wollmann, 2018). My work extends this literature by introducing the business cycle into these dynamics and evaluating their welfare implications in a broader context. I am not the first to incorporate aggregate fluctuations into the study of industry dynamics (Collard-Wexler, 2013; Jeon, 2022); however, I build on this line of research by explicitly accounting for producers' entry and exit decisions alongside heterogeneous consumer demand.

Outline. This paper is structured as follows. Section 2 describes the data and presents motivating facts about the U.S. product market over the period surrounding the Great Recession. Section 3 introduces the model of product market dynamics. Section 4 outlines the estimation methods. Section 5 presents results derived from product-level microdata. Section 6 presents counterfactual analyses that quantify the welfare costs of the business cycle. Finally, Section 7 concludes the paper.

2 Motivating Facts

This section introduces the data and presents motivating facts that illustrate how business cycles generate heterogeneous impacts across product markets. Section 2.1 describes the micro-level data from the U.S. consumer goods sector used in the analysis. Section 2.2 documents stylized facts that highlight differences in macroeconomic dynamics at both the aggregate and disaggregated levels.

2.1 Data Description

This paper uses the NielsenIQ Retail Scanner dataset from the Kilts Center for Marketing to study how the business cycle shapes product market dynamics. The dataset records prices and quantities sold at the Universal Product Code (UPC), store, and week levels, covering billions of transactions from 30,000–50,000 retail stores across the United States. It captures more than half of total sales in grocery and drug stores and about 30% of sales in mass merchandisers nationwide, providing a representative view of U.S. consumption.[6] Because of its scale and granularity, the dataset has become widely used in macroeconomic research on business cycles, firm dynamics, inflation, fiscal policy, and household inequality, among other topics.[7]

The dataset begins in 2006, so this study focuses on the business cycle episode surrounding the Great Recession. Focusing on this single recessionary episode still provides valuable evidence, as the Great Recession was one of the most severe demand-driven recession—well captured by the current framework based on aggregate demand shocks. While the analysis could, in principle, be extended to the COVID-19 recession, identifying producers' dynamic decisions during that period is complicated by multiple confounding factors. Accordingly, the analysis uses data from 2006 to 2019 and follows the definitions of entry and exit established in previous studies (Broda and Weinstein, 2010; Argente, Lee, and Moreira, 2018; Granja and Moreira, 2023). Entry is defined as the first period in which a product is sold in any store after at least four consecutive quarters with no prior sales, and exit as the last period in which the product is sold, followed by four consecutive quarters with no transactions. This timing rule improves the reliability of identifying entry and exit events by minimizing the influence of seasonality and naturally captures extensive-margin decisions at the national level over time.

Consistent with these definitions, the dynamic model estimation sample begins in the fourth quarter of 2007 (the onset of the Great Recession) while pooled data from 2006 through 2019 are used for demand estimation. To balance computational tractability with the need to capture structural changes over the business cycle, I divide the dynamic estimation period into three phases: 2007Q4–2009Q2, 2009Q3–2012Q4, and 2013Q1–2016Q4. Observations prior to 2007Q4 and after 2016Q4 are treated as normal periods without major aggregate demand shocks. All

[6] See https://www.chicagobooth.edu/research/kilts/research-data/nielseniq for details.
[7] For a review of studies using scanner data, see Dubois, Griffith, and O'Connell (2022).

estimations are conducted separately by industry, defined as the set of products within NielsenIQ categories, under the assumption that producers' decisions are independent across industries and that consumers' choices do not consider products from other categories.[8]

Throughout the paper, products are defined at the distinct brand level to capture firms' strategies of product differentiation, following Nevo (2001). Defining products at the brand level reflects the idea that differentiation in consumer goods markets primarily occurs across brands rather than among individual UPCs within a brand. Each UPC at the barcode level is consolidated into a standardized unit measure (e.g., ounces, liters) within each brand, accounting for variations in package sizes and multipack configurations.[9] Prices and quantities for each product are therefore expressed as price per unit and total units sold. This aggregation distinguishes the present study from previous research that examined cyclical properties of product markets at the UPC level. All prices are deflated using the Core Consumer Price Index.

I further aggregate the data to higher levels to facilitate a more practical analysis, consistent with previous research (Nevo, 2001; Döpper et al., 2025). Weekly observations are aggregated to a quarterly frequency, which is standard in the business cycle literature. This quarterly aggregation mitigates potential biases from consumers' dynamic demand behavior and reflects firms' realistic decision-making horizon for product entry and exit.[10] Furthermore, sales data from individual stores are aggregated at the retail-chain level, so the uniform-pricing property across national retail chains does not affect the analysis (DellaVigna and Gentzkow, 2019). Regional aggregation is also performed at Nielsen's Designated Market Areas (DMAs), which group together counties that share local television advertising markets.[11] Thus, throughout the paper, market coverage is defined at the chain and DMA-region level rather than at the disaggregated physical-store level. For the demand estimation, I use price and quantity information at this local market level. The intensive-margin decision (market coverage) is measured consistently across both the disaggregated local and aggregated national levels.

[8]For instance, categories such as soft drinks–carbonated" and soft drinks–low calorie" are not combined, as NielsenIQ does not group them together. This assumption simplifies estimation but does not capture potential substitution across categories.

[9]Appendix Figure A1 shows that most brands consist of only a few unique UPCs, suggesting that the assumption of a single product at the brand level is not restrictive. See Appendix A for additional details.

[10]For example, consumers are unlikely to stockpile cereals three months in advance, and producers do not adjust their entry or exit decisions on a weekly basis.

[11]See https://www.nielsen.com/dma-regions/.

Figure 1: Aggregate Number of Incumbent Products over Time

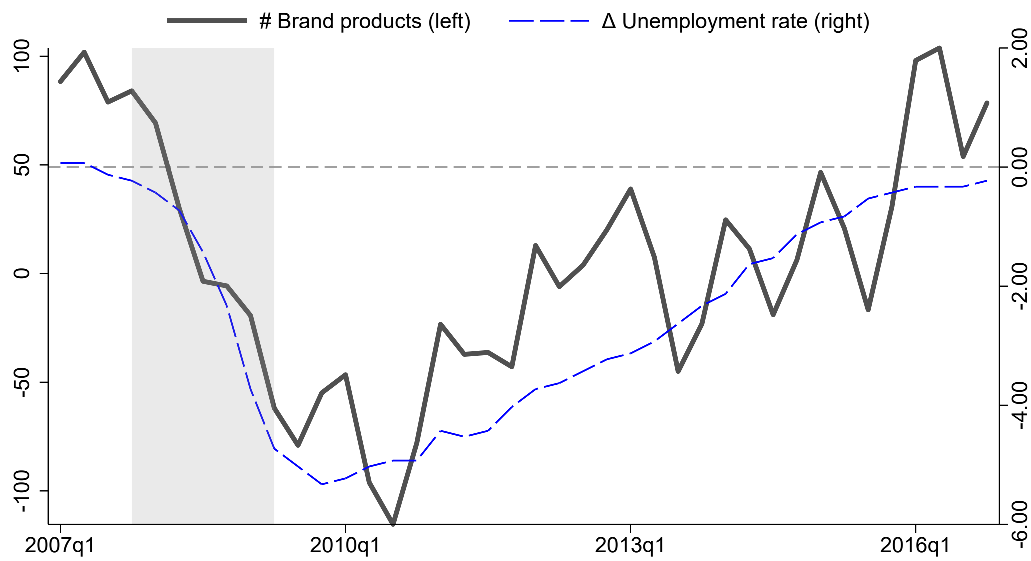

Notes: This figure plots the time series of detrended incumbent products and normalized unemployment rates. The solid black line (left axis) shows the number of incumbent products across industries in the NielsenIQ sample used for analysis. The dashed blue line (right axis) represents the aggregate economy, measured by the normalized unemployment rate, defined as the deviation from its pre-recession mean. The number of products is detrended using the Hodrick–Prescott filter after applying a three-quarter moving average. The shaded area indicates the NBER-defined recession period.

2.2 Heterogeneous Product Market Dynamics

This paper is motivated by the fact that product market dynamics vary substantially across industries, even when aggregate patterns appear similar. Figure 1 illustrates how such aggregation can mask important heterogeneity. The number of incumbent product combinations across industries in the sample closely comoves with the normalized unemployment rate, dropping sharply during the Great Recession and recovering at a similar pace as the aggregate economy. The correlation between the two series is striking (0.85). This procyclical behavior is consistent with prior evidence that product entry and exit are highly sensitive to macroeconomic conditions (Broda and Weinstein, 2010; Argente, Lee, and Moreira, 2018), as fluctuations in the total number of products are driven by changes in entry and exit dynamics.

Examining these dynamics at the industry level reveals substantial heterogeneity, suggesting that the economic consequences of aggregate shocks may diverge across product space. Figure 2 plots this variation by showing the nine product categories with the largest sales shares in the sample. The black lines illustrate the number of incumbent products in each industry, normalized

Figure 2: Number of Incumbent Products across Industries

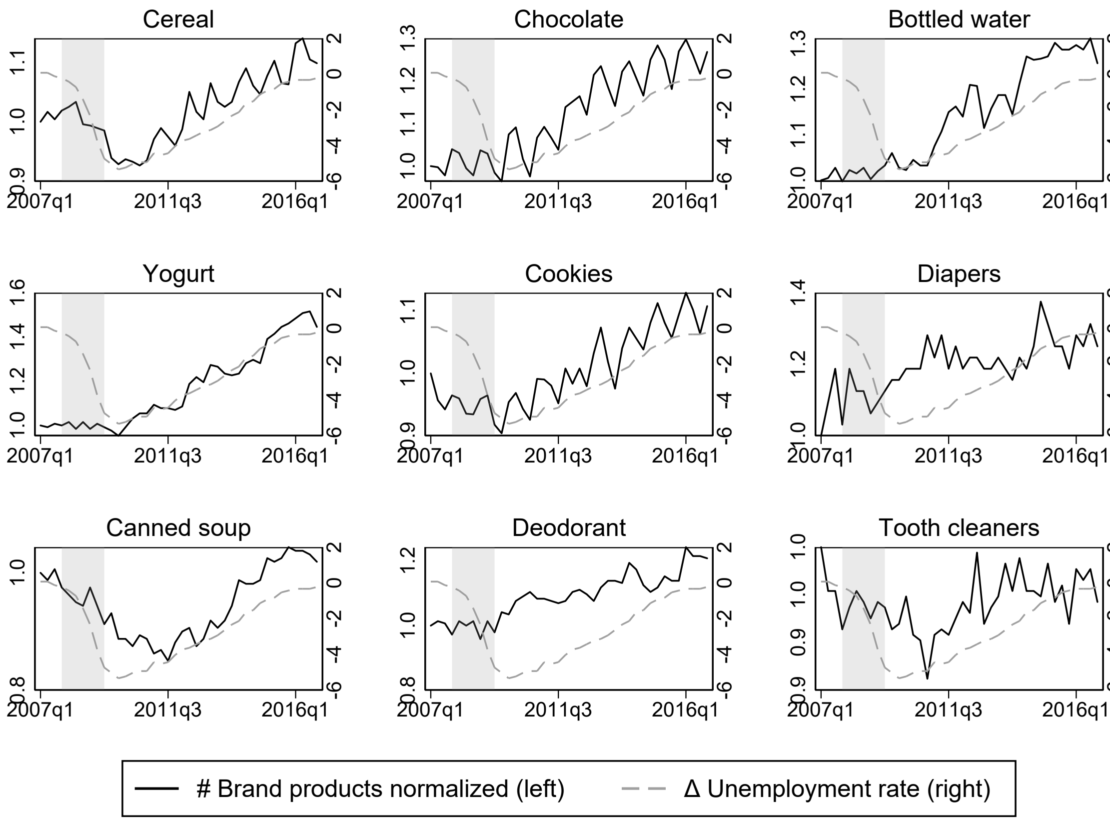

Notes: This figure plots the time series of the number of incumbent products (solid black line, left axis) for each industry, normalized to its 2007Q1 level. The nine panels display the top nine product categories ranked by average revenue in the analysis sample. The dashed gray line (right axis) shows the normalized unemployment rate. Shaded areas indicate NBER-defined recession periods.

to their 2007Q1 levels. The figure shows that some industries exhibit strong procyclical patterns, whereas others appear largely unaffected. Moreover, the magnitude and timing of fluctuations vary. For instance, industries such as cereal, canned soup, and tooth cleaners experience sharp declines in product variety during the recession, whereas chocolate, yogurt, and deodorants display relatively stable trends. These differences indicate that recessions do not uniformly affect product markets, underscoring the importance of understanding the mechanisms behind this heterogeneity.

Finally, Figure 3 further highlights this heterogeneity by comparing the impact of aggregate shocks across industries. It reports the change in the average number of products between the pre-recession and post-recession periods for the nine industries that experienced the largest declines in the sample. For example, pre-shave cosmetics (such as shaving cream) experienced a decline of

Figure 3: Changes in the Average Number of Product between Pre-recession and Post-recession

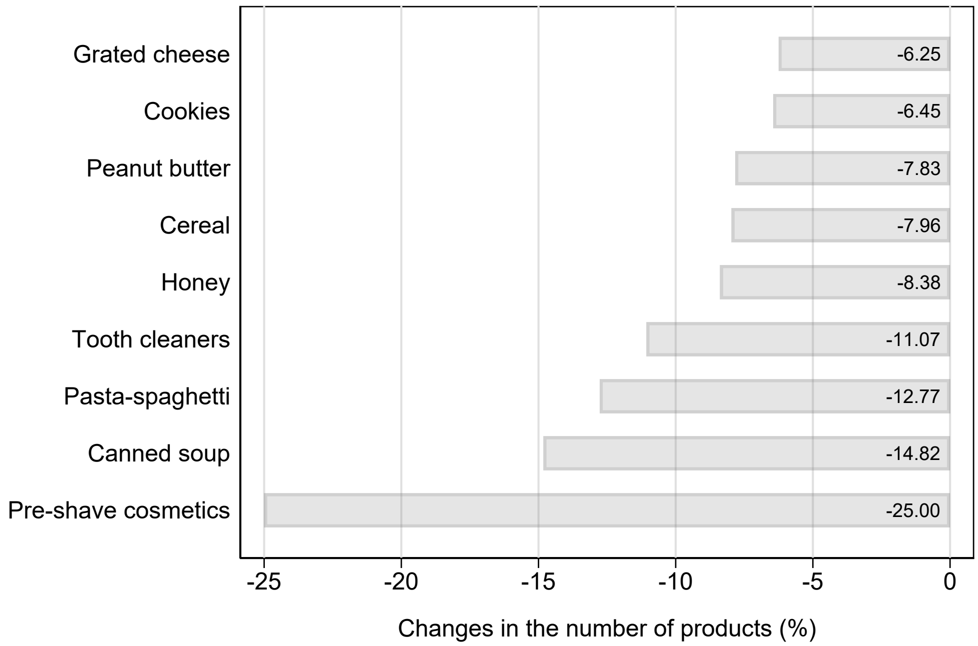

Notes: This figure plots the percentage change in the average number of incumbent products after the recession relative to the pre-recession period for each industry. The figure displays the nine categories that experienced the largest declines.

about 25% in the number of products, whereas cookies fell by around 6.5%. Some of the industries that appeared strongly affected in the previous figure are not included here. Thus, comparing only pre- and post-recession averages would overlook substantial within-period variation in product market dynamics, underscoring the importance of capturing these dynamics in the analysis.

This paper argues that the heterogeneous responses of product markets to aggregate demand shocks arise from differences in demand structures within and across industries. Explaining them requires recovering the underlying demand parameters. Moreover, to fully capture product market dynamics, it is essential to account for producers' dynamic decisions: entry, exit, and adjustments in market coverage. Throughout the analysis, I show how heterogeneous demand structures shape these dynamics and, in turn, their macroeconomic implications.

3 Model

This section develops a dynamic model of the product market, following the moment-based Markov equilibrium (MME) framework of Ifrach and Weintraub (2017) in the spirit of the dynamic oligopoly model of Ericson and Pakes (1995).[12] It also incorporates aggregate demand shocks to capture business cycle fluctuations. The model further includes a discrete-choice demand system with heterogeneous consumers (Berry, Levinsohn, and Pakes, 1995). The subsections describe the model environment (Section 3.1), producers' moment-based strategies (Section 3.2), and their dynamic optimization problems (Section 3.3).

3.1 Environment

Time. Time is discrete and finite. A period corresponds to a quarter, indexed by $t = 1, 2, \ldots, \bar{T}$.

Producers. In each period, there is a finite number of incumbent producers, denoted J_t. Each producer indexed by j is a single-product firm, so that product characteristics are sufficient to describe the producer.[13] In addition, there are J_t^{pe} potential entrants, representing the maximum number of producers the market can accommodate.

Producer state. Incumbent producer heterogeneity is characterized by two key state variables: *product appeal*, ξ_{jt}, and *market coverage*, N_{jt}. In each period, the individual state of producer j is denoted by the pair $x_{jt} = (\xi_{jt}, N_{jt}) \in \mathcal{X}$, where $\mathcal{X}$ denotes the finite set of possible state combinations. Product appeal captures favorable shifts in demand conditional on price and other observed characteristics. It follows a first-order Markov process:

$$\xi_{jt+1} = \mathbb{E}[\xi_{jt+1} \mid \xi_{jt}] + \eta_{jt+1}, \tag{1}$$

where η_{jt+1} is an idiosyncratic demand shock drawn from an i.i.d. standard normal distribution. Market coverage reflects a product's availability to consumers, measured by the number of local retail stores in which it is sold.[14] It is determined by producers each period and evolves over time through costly endogenous adjustment decisions. Potential entrants are ex ante homogeneous; their heterogeneity is realized upon entry as new producers.

[12] The MME framework has been widely applied in recent studies, for example to banking industry dynamics (Corbae and D'Erasmo, 2021) and supply chain management through outsourcing (Thurk, 2024).

[13] This simplification is made for practical reasons, as estimating entry and exit dynamics for heterogeneous products within multi-product firms is computationally infeasible and highly challenging (Benkard, 2004).

[14] The total number of stores in the economy is fixed over time.

Aggregate state. The aggregate state of the market is characterized by z_t and $\hat{s}_t$. The aggregate demand shock $z_t \in \mathcal{Z}$ represents the current level of market demand common to all producers. It captures the mass of consumers actively purchasing in the market; adverse shocks reduce this mass, shrinking the effective market size faced by producers. The aggregate demand shock follows a first-order Markov process:

$$z_{t+1} = \mathbb{E}[z_{t+1} \mid z_t] + \zeta_{t+1}, \tag{2}$$

where ζ_{t+1} is a white-noise shock. The variable $\hat{s}_t \in \mathcal{S}$ denotes the aggregate utility index, which summarizes the attractiveness of all products available in period t. Producers condition on this aggregate moment because it provides information about their product j's profitability relative to industry-level demand. A detailed definition of $\hat{s}_t$ is provided in Section 3.2.

Exit process. In each period, incumbent producer j decides whether to remain in the market or exit. The exit decision is represented by the action a_{jt}, where $a_{jt} = 0$ indicates exit. If the producer exits, they receive a private exit value γ_{jt}^{ex} drawn from an exponential distribution with mean γ^{ex}.

Adjustment process. If the producer decides to stay, they choose the level of next period's market coverage, N_{jt+1}. The action of a staying producer is denoted by $a_{jt} = N_{jt+1}$ and evolves according to the following endogenous adjustment process:

$$N_{jt+1} = N_{jt} + \Delta_N, \tag{3}$$

where Δ_N denotes the change in market coverage. The realization of next period's market coverage is subject to an exogenous noise term ν_{jt}, drawn from an i.i.d. Type I extreme value distribution with scale parameter ψ. Adjusting market coverage away from its current level incurs an adjustment cost.

Entry process. In each period, each potential entrant decides whether to enter the market in the next period. The decision is denoted by a_{jt}^{pe}, where $a_{jt}^{pe} = 1$ indicates entry and $a_{jt}^{pe} = 0$ otherwise. If the entrant decides to enter, they incur a one-time private entry cost γ_{jt}^{en} drawn from an exponential distribution with mean γ^{en}. After paying the entry cost, the entrant's initial state is realized. All new entrants start with the lowest level of market coverage, $N_{jt+1}^{en} = 1$, and their product appeal ξ_{jt+1}^{en} is drawn from a time-invariant distribution G.

Consumer problem. In each period t, there are M_t consumers, indexed by i. Consumers are myopic and maximize static utility by choosing the product j that provides the highest utility, given price p_{jt}, market coverage N_{jt}, and product appeal ξ_{jt}. Each consumer is heterogeneous in price sensitivity, denoted by α_i, which follows a normal distribution with mean α and variance σ_α^2. In addition, consumer i's idiosyncratic taste for product j, denoted by ϵ_{ijt}, follows an i.i.d. Type I extreme value distribution.

If consumer i chooses product j among the J_t available products, the choice-specific indirect utility is given by

$$v_{ijt} = \alpha_i p_{jt} + \phi \log N_{jt} + \xi_{jt} + \epsilon_{ijt}. \tag{4}$$

The probability that consumer i chooses product j equals the probability that product j yields the highest utility among all available products, including the outside option. This probability defines consumer i's demand for product j, denoted d_{ijt}. Using the statistical property of ϵ_{ijt} under the Type I extreme value assumption, we can express d_{ijt} as

$$d_{ijt} = \frac{\exp[\alpha_i p_{jt} + \phi \log N_{jt} + \xi_{jt}]}{1 + \sum_{k \in J_t} \exp[\alpha_i p_{kt} + \phi \log N_{kt} + \xi_{kt}]}, \tag{5}$$

where the utility of the outside option is normalized to zero.

3.2 Moment-based strategies

Demand for product j depends not only on its own state but also on the states of all other producers. Equation (5) illustrates this interdependence: d_{ijt} is a function not only of $x_{jt} = (\xi_{jt}, N_{jt})$ but also of $x_{kt} = (\xi_{kt}, N_{kt})$ for $k \neq j$. However, it is difficult for producers to track the states of all competitors over time, as they evolve each period. Thus, producers summarize this information using an aggregate moment that serves as a sufficient statistic for the distribution of competitors' states rather than tracking individual competitors' states over time (Ifrach and Weintraub, 2017).

The aggregate moment in this framework is the aggregate utility index $\hat{s}_t$, defined as

$$\hat{s}_t = \sum_{x \in \mathcal{X}} \mu_t(x) \exp\left[\alpha_i p_t(x) + \phi \log N_t(x) + \xi_t(x)\right]. \tag{6}$$

Let μ_t denote the distribution of producers across states. Then $\mu_t(x)$ represents the number of producers in state $x \in \mathcal{X}$, while $\xi_t(x)$ and $N_t(x)$ denote the levels of product appeal and market coverage associated with that state, respectively. This index summarizes the overall valuation of all products in the market and, hence, affects each producer's own demand—for instance, when this valuation rises, other things equal, the relative utility of product j declines. For each producer j, this aggregate moment provides a key signal about its own profitability, on which producers rely when making dynamic decisions.

Now, I define the moment-based strategies for each producer as functions of the moment-based state $\hat{s}_t$, consistent with the assumption that strategies depend on a few sufficient statistics of the industry state rather than the entire distribution. Let λ_t denote the profile of all producers' actions, $\lambda_t = \{a_{jt}\}_{j \in J_t} \cup \{a_{jt}^{pe}\}_{j \in J_t^{pe}}$, which depends on individual states, the aggregate state, and idiosyncratic shocks. Given λ_t, which summarizes all producers' strategies at time t, the distribution of producers μ_t evolves according to

$$\mu_{t+1}(x') = \begin{cases} \sum_{x \in \mathcal{X}} Q_t(x' \mid x, \lambda_t)\mu_t(x) + \left(\sum_{j \in J_t^{pe}} a_{jt}^{pe}\right) dG(\xi') & \text{if } N' = N^{en}, \\ \sum_{x \in \mathcal{X}} Q_t(x' \mid x, \lambda_t)\mu_t(x) & \text{otherwise}, \end{cases} \tag{7}$$

where $x' = (\xi', N')$ and $Q_t(x' \mid x, \lambda_t)$ denotes the transition kernel governing the evolution of incumbents' states conditional on the strategy profile. In what follows, I assume that the aggregate utility index $\hat{s}_t$ evolves approximately as a Markov process.

3.3 Dynamic Optimization

At the beginning of period t, the individual states of all incumbent producers, $x_{jt} = (\xi_{jt}, N_{jt})$, the aggregate demand shock z_t, and the aggregate industry state $\hat{s}_t$ are realized. Individual demand can be expressed as $d_{ijt}(\xi_{jt}, N_{jt}, \hat{s}_t; \alpha, \sigma_\alpha^2, \phi)$, which depends on the producer's own state, the market state, and the demand parameters. Combining the individual and aggregate components, I denote the overall state at time t as $s_{jt} = (\xi_{jt}, N_{jt}, \hat{s}_t, z_t)$ for notational simplicity. Noting that the aggregate demand shock z_t scales the overall market size, I obtain the aggregate quantity demanded for product j by integrating over the distribution of consumer price sensitivities:

$$D_{jt}(s_{jt}; \alpha, \sigma_\alpha^2, \phi) = \left(\int d_{ijt}(\xi_{jt}, N_{jt}, \hat{s}_t; \alpha, \sigma_\alpha^2, \phi) \, dF(\alpha_i) \right) M_t \exp(z_t). \tag{8}$$

In each period, an incumbent producer earns static profit

$$\pi_{jt}(s_{jt};\alpha,\sigma_\alpha^2,\phi) = \left[p_{jt}^*(\xi_{jt},N_{jt}) - mc_{jt}(s_{jt};\alpha,\sigma_\alpha^2,\phi)\right] D_{jt}(s_{jt};\alpha,\sigma_\alpha^2,\phi), \tag{9}$$

where $p_{jt}^*(\xi_{jt},N_{jt})$ is the observed equilibrium price expressed as a function of individual states that evolve over time, and $mc_{jt}(s_{jt};\alpha,\sigma_\alpha^2,\phi)$ is the marginal cost. To maintain tractability, I do not explicitly model pricing decisions within the dynamic framework. Instead, this model assumes that producers compete in prices according to a static Bertrand–Nash pricing game in each period t, and that the corresponding first-order condition holds. Within the dynamic model, the marginal cost $mc_{jt}(s_{jt};\alpha,\sigma_\alpha^2,\phi)$ is recovered from this static first-order condition using observed prices and the estimated demand system.[15]

After earning period profit, incumbents observe their private shock components related to exit, γ_{jt}^{ex}, and market coverage adjustment, ν_{jt}. Each incumbent then decides whether to stay in the market or exit. In making this decision, they consider the expected continuation value, discounted by $\beta \in (0,1)$, and the action profile λ_t that summarizes other producers' strategies. The dynamic optimization problem of an incumbent is summarized by the following Bellman equation:

$$\begin{aligned} V_t(s_{jt},\lambda_t;\theta) = \pi_t(s_{jt};\alpha,\sigma_\alpha^2,\phi) + \mathbb{E}_{\gamma^{ex}}\Bigg[\max\Bigg\{ & \gamma_{jt}^{ex}, \\ & \mathbb{E}_\nu\left[\max_{N'}\left\{ -C(N',N_{jt};\gamma^c) + \nu_{jt} + \beta\mathbb{E}_t\left[V_{t+1}(s_{jt+1},\lambda_{t+1};\theta)\middle| s_{jt},\lambda_t\right]\right\}\middle| s_{jt}\right]\Bigg\}\Bigg| s_{jt}\Bigg], \end{aligned} \tag{10}$$

where $\theta = \{\alpha,\sigma_\alpha^2,\phi,\gamma^{en},\gamma^{ex},\gamma^c,\psi\}$, and $\gamma^c = (\gamma_0,\gamma_1^+,\gamma_1^-,\gamma_2)$. The adjustment cost function is defined as

$$\begin{aligned} C(N_{jt+1},N_{jt};\gamma^c) =& \mathbb{1}\{N_{jt+1} > N_{jt}\}\left(\gamma_0 + \gamma_1^+\Delta_N + \gamma_2\Delta_N^2\right) \\ &+ \mathbb{1}\{N_{jt+1} < N_{jt}\}\left(\gamma_0 + \gamma_1^-\Delta_N + \gamma_2\Delta_N^2\right), \end{aligned} \tag{11}$$

where the linear adjustment parameters are asymmetrically specified.

At the same time, each potential entrant draws a private entry cost γ_{jt}^{en}. Entrants decide whether to enter the market by comparing this cost with their expected discounted value upon entry, given

[15] Under the single-product assumption, the recovered marginal costs may be biased for multi-product firms.

the current aggregate states and action profile, by solving the following problem:

$$\mathbb{E}_{\gamma^{en}}\left[\max\left\{0,\ \beta\mathbb{E}_t\Big[V_{t+1}\left(s_{jt+1},\lambda_{t+1};\theta\right)\big|\hat{s}_t,z_t,\lambda_t\Big]-\gamma_{jt}^{en}\right\}\Bigg|\hat{s}_t,z_t\right]. \tag{12}$$

Producers that choose to enter become incumbents in period $t+1$, with initial states determined by a product-appeal draw $\xi_{jt+1}\sim G(\cdot)$ and $N_{jt}^{en}=1$. From period $t+1$ onward, they solve the dynamic optimization problem as incumbents.

This paper adopts the moment-based Markov equilibrium framework of Ifrach and Weintraub (2017). To ensure the existence of an equilibrium and computational feasibility, I assume that the industry's nonstationary evolution converges to its stationary counterpart over a sufficiently long time horizon (Benkard, Jeziorski, and Weintraub, 2024). The dynamic problem is solved by backward induction, where the terminal period's value function is approximated as if no producers exit the market and each continues to earn profits indefinitely without adjusting market coverage, while only the product appeal ξ_{jt} evolves according to its transition probability.

In equilibrium, the moment-based strategy λ_t satisfies the following conditions for all periods t: incumbent producers solve (10), potential entrants solve (12), aggregate demand shocks evolve according to (2), the number of producers evolves according to (7), and each producer forms beliefs about the industry state based on the moment-based states in (6).

4 Estimation

This section describes a two-step estimation procedure. In Section 4.1, I estimate key demand parameters in a static framework following Berry, Levinsohn, and Pakes (1995), approximating the demand equation using the method of Salanié and Wolak (2022) to handle multiple industries efficiently. In Section 4.2, I estimate the parameters governing producers' dynamic optimization by maximum likelihood using a nested fixed-point algorithm (Rust, 1987; Igami, 2017).

4.1 Demand Estimation

The demand structure of this paper is described in equation (4). Consumers are heterogeneous in their price sensitivity (Berry, Levinsohn, and Pakes, 1995). Each industry has its own demand parameters for the price sensitivity distribution $(\alpha, \sigma_\alpha^2)$ and the preference for product availability ϕ. To capture heterogeneity both within and across industries, I estimate the demand parameters separately for each industry. However, estimating random-coefficient logit demand models is computationally intensive even for a single industry, and extending this approach to many industries substantially increases the computational burden. To address this challenge, I follow the approach of Salanié and Wolak (2022), which approximates the nonlinear demand system with a linear equation. This method has also been applied to multiple product categories in recent work by Brand (2021), who show that results from both methods are qualitatively similar.

For the demand estimation, I use information on prices and quantities at the local market level, defined as the combination of region and retail chain in each quarter, rather than at the national level. This unit of analysis differs from the dynamic estimation, where I assume that entry and exit decisions are made at the national level on a quarterly basis. I assume that consumer preference parameters within each industry are identical and remain unchanged across levels of aggregation. The rationale is as follows. First, this unit of estimation is standard in the industrial organization literature, as it exploits variation in demand across local markets. Second, obtaining consistent estimates of demand parameters under this approximation requires a sufficiently large number of markets (Salanié and Wolak, 2022). For these reasons, I conduct the demand estimation at this level to ensure that my estimated parameters are consistent with findings in the literature. Finally, to maximize the number of markets, I pool data across all quarters from 2006 to 2019 within each industry for estimation.

In this approximated estimation, the demand system can be estimated in a manner similar to the standard logit demand model. Following Berry (1994), we apply the logit inversion using the outside-option market share to equation (8). Since I estimate the demand system at the local market level, each market ℓ's aggregate demand is expressed for this inversion. Applying the approximation method, the logit inversion can be expressed as

$$\log\left(\frac{D_{j\ell t}}{D_{0\ell t}}\right) = \alpha p_{j\ell t} + \sigma_\alpha^2 K_{j\ell t} + \phi \log N_{j\ell t} + \xi_{j\ell t} + \mathcal{O}(\cdot), \tag{13}$$

where $K_{j\ell t}$ is an *artificial regressor* defined as

$$K_{j\ell t} = \left(\frac{p_{j\ell t}}{2} - e_{\ell t}\right) p_{j\ell t}, \tag{14}$$

and

$$e_{\ell t} = \sum_{j \in J_{\ell t}} D_{j\ell t} p_{j\ell t}, \tag{15}$$

which represents the prices in market ℓ weighted by market shares. For each product j, the market coverage $N_{j\ell t}$ is the same across markets; thus, $N_{j\ell t} = N_{jt}$.

Estimating equation (13) requires two instrumental variables (one for price and one for the artificial regressor, which contains both market shares and prices). I employ three instruments: a quadratic version of the differentiation instruments proposed by Gandhi and Houde (2019), constructed using prices and the number of stores (at the chain $\times$ region level) in which product j is sold, and the local unemployment rate (Backus, Conlon, and Sinkinson, 2021).[16] To construct the price instrument, I first create the Hausman (1996) instrument and then extract retailer-specific price changes by subtracting the national average, following Allcott, Lockwood, and Taubinsky (2019) and Conlon and Gortmaker (2023). I then recover predicted prices by regressing observed prices on the Hausman instrument with a set of fixed effects; these predicted prices are subsequently used to generate the differentiation IV.

For the empirical specification, I include product fixed effects and market fixed effects. The aggregate demand shocks are absorbed by these fixed effects, so the recovered demand parameters should be interpreted as conditional on these shocks. To compute market shares, I define the market size for each local market as the total population multiplied by the number of actual stores, measured at the disaggregated store level. I also normalize market shares so that the combined share of inside goods is approximately 0.45, following Döpper et al. (2025). The estimation sample for the demand system includes products with private labels. Because it is not possible to identify the manufacturers of these products, I treat their dynamics as exogenous and do not model their endogenous decisions; they are therefore excluded from the dynamic estimation sample. After estimation, I recover the product appeal $\xi_{j\ell t}$ by adding the recovered product fixed effect and then

[16] A chain is defined at the parent-company level, and a region follows Nielsen's Designated Market Areas (DMAs).

Table 1: Estimates of Demand Parameters

Parameters	Notation	Cereal	Yogurt
Price sensitivity (mean)	α	–14.95	–38.83
		(3.03)	(5.46)
Price sensitivity (variance)	σ_α^2	15.71	71.79
		(5.99)	(37.16)
Product availability	ϕ	0.91	0.93
		(0.00)	(0.05)
Autocorrelation of product appeal		0.83	0.80
Median price elasticity of demand		–2.57	–4.50

Notes: Robust standard errors are in parentheses. Results for two industries commonly studied in the literature are presented as illustrative examples. The product-level price elasticity of demand is computed as $\frac{\partial D_{jt}}{\partial p_{jt}} \cdot \frac{p_{jt}}{D_{jt}}$. Median price elasticities are calculated within each industry as the median across all sample periods.

aggregate both product appeal and prices at the national level using quantity weights.

Table 1 reports the estimated demand parameters and the recovered median price elasticities for the cereal and yogurt industries as illustrative examples. The current estimation sample covers 15 industries in total. The estimates align well with values documented in the literature, with the median price elasticity of ready-to-eat cereal calculated at –2.57 and yogurt at –4.50.[17] These estimates fall within the range reported in previous studies: for cereals, from –2.3 to –3.5 (Nevo, 2001; Backus, Conlon, and Sinkinson, 2021; Atalay et al., 2023; Döpper et al., 2025), and for yogurt, from –3.1 to –5.0 (mean or median) (Hristakeva, 2022; Atalay et al., 2023; Döpper et al., 2025).

The results highlight two key points. First, there is notable within-industry heterogeneity, as captured by the positive and significant variance parameters. Consumers dislike higher prices, but the degree of this price sensitivity varies across individuals. Second, the demand structures differ significantly between the two industries. The mean price sensitivity in yogurt is more than twice that of cereal, and the variance of price sensitivity is also substantially larger. This indicates meaningful across-industry heterogeneity. In the following analysis, I examine how these differences in demand structures contribute to the heterogeneous responses of product market dynamics to aggregate shocks. Estimation results for other industries are presented in Appendix F.

[17] Integration is performed using Gauss–Hermite quadrature of order 11.

4.2 Estimating Dynamic Model

This paper estimates the dynamic parameters using a full-solution approach with a nested fixed-point algorithm, following Rust (1987) and Igami (2017). I use maximum likelihood estimation to rationalize the observed product market behavior. For the state-space specification, product appeal is discretized into 11 grid points, and transition probabilities are computed directly from the data using Laplace smoothing. The market coverage is discretized into 20 grid points. Unlike product appeal—which is recovered as a residualized variable from the demand estimation—market coverage is right-skewed with a long tail. To reduce the state-space dimension, I cap the maximum at 200 and set the grid size to 10, starting from 1.[18]

Because product market dynamics during the Great Recession are highly nonlinear and vary substantially over time, estimating the model over the entire sample period (2007–2018) would not adequately capture these transitional patterns. The environment is nonstationary: aggregate demand and producers' decisions evolve over time across distinct phases rather than remaining constant. Therefore, I estimate the model separately for three subperiods: 2007Q4–2009Q2, 2009Q3–2012Q4, and 2013Q1–2016Q4. For the pre-recession period and after 2016, I assume that the economy returns to normal conditions without aggregate demand shocks.[19]

The dynamic estimation recovers the parameters that govern producers' entry and exit decisions, $(\gamma^{\text{en}}, \gamma^{\text{ex}})$, and market coverage adjustment decisions, $(\gamma_0, \gamma_1^+, \gamma_1^-, \gamma_2)$. The objective is to identify the parameter values that best rationalize the observed evolution of each product over time. This is accomplished by maximizing the likelihood of the observed data conditional on the model-implied optimal policies. One exception is the logit scale parameter ψ, which governs the stochasticity of market coverage adjustments. This parameter is estimated separately using a minimum-distance procedure that matches the model-implied aggregate number of products over time to its empirical counterpart, after recovering the remaining dynamic parameters.

[18] In the pooled sample across all industries, 95 percent of observations involve fewer than 100 stores at the chain and region level. For a detailed discussion and summary statistics, see Appendix B.

[19] In the current framework, I model the aggregate demand shock as following a deterministic trajectory proxied by the normalized unemployment rate to capture the time-varying environment during and after the Great Recession. This approach rationalizes observed transitions in product market behavior while maintaining computational tractability. However, it abstracts from stochastic aggregate shocks that could alter firms' expectations and reoptimization decisions each period. Benkard, Jeziorski, and Weintraub (2024) develop a framework that explicitly incorporates such stochastic aggregate shocks within a nonstationary Markov equilibrium, showing that equilibrium strategies may deviate when aggregate uncertainty is high. Extending the current model in this direction is an important avenue for future research.

I begin by defining the producer's policy functions, which can be expressed in closed form. The incumbent's dynamic optimization problem is summarized in Equation (10). For simplicity, and with a slight abuse of notation, I omit model parameters from the notation. The choice-specific value function for incumbent j in state s_{jt}, when selecting next-period market coverage N^k, is defined as

$$v_{jt}^{k} = -C(N^k, N_{jt}) + \beta\,\mathbb{E}_t[V_{t+1}(s_{jt+1}, \lambda_{t+1}) \mid s_{jt}, \lambda_t]\,, \tag{16}$$

where k indexes the discrete level of market coverage N.

Given that the adjustment of market coverage is subject to an idiosyncratic shock ν_{jt}, which follows a Type I extreme value distribution with scale parameter ψ, the expected ex-ante continuation value conditional on staying in the market can be written as

$$\tilde{V}_{jt} = \psi \left[\gamma_E + \log \sum_k \exp\left(\frac{v_{jt}^k}{\psi} \right) \right], \tag{17}$$

where γ_E denotes the Euler constant (≈ 0.57721).

Producers exit the market when their privately observed exit shock γ_{jt}^{ex} exceeds the expected continuation value $\tilde{V}_{jt}$, where γ_{jt}^{ex} follows an exponential distribution with mean parameter $\gamma^{\text{ex}} > 0$. Using the properties of the exponential distribution, the probability that incumbent j exits the market, conditional on state s_{jt} and action profile λ_t, is given by

$$\Pr_t(a_{jt} = 0 \mid s_{jt}, \lambda_t) = \begin{cases} \exp\left(-\dfrac{\tilde{V}_{jt}}{\gamma^{\text{ex}}} \right), & \tilde{V}_{jt} \geq 0, \\ 0, & \tilde{V}_{jt} < 0. \end{cases} \tag{18}$$

Similarly, the probability of staying in the market and selecting next-period market coverage level $N_{j,t+1}^k$ among all possible levels n can be expressed as

$$\Pr_t(a_{jt} = N_{j,t+1}^k \mid a_{jt} > 0, s_{jt}, \lambda_t) = \Big(1 - \Pr_t(a_{jt} = 0 \mid s_{jt}, \lambda_t) \Big) \frac{\exp\left(\dfrac{v_{jt}^k}{\psi} \right)}{\sum_n \exp\left(\dfrac{v_{jt}^n}{\psi} \right)}, \tag{19}$$

where the first term, $1 - \Pr_t(a_{jt} = 0 \mid s_{jt}, \lambda_t)$, denotes the probability of remaining in the market, and the second term represents the logit choice probability over the discrete set of market coverage levels, conditional on survival.

Using the above expressions, the memoryless property of the exponential distribution, and the per-period profit function in Equation (9), the incumbent's Bellman equation can be expressed in closed form as

$$V_t(s_{jt}, \lambda_t) = \pi_{jt}(s_{jt}) + \Pr_t(a_{jt} = 0 \mid s_{jt}, \lambda_t)\big(\gamma^{\text{ex}} + \tilde{V}_{jt}\big) + \big(1 - \Pr_t(a_{jt} = 0 \mid s_{jt}, \lambda_t)\big)\tilde{V}_{jt}. \tag{20}$$

We also define the probability of entry for potential entrants. Each potential entrant draws an idiosyncratic entry cost γ_{jt}^{en} from an exponential distribution with mean parameter $\gamma^{\text{en}} > 0$. The probability that a potential entrant enters the market is given by

$$\Pr_t(a_{jt}^{\text{pe}} = 1 \mid s_{jt}, \lambda_t) = 1 - \exp\left(-\frac{\beta\, \mathbb{E}_t[V_{t+1}(s_{jt+1}, \lambda_{t+1}) \mid s_{jt}, \lambda_t]}{\gamma^{\text{en}}}\right), \tag{21}$$

and it equals zero whenever the expected value of entry is negative.

The conditional choice probabilities solved from the model for incumbents are denoted by $\Pr_t(a_{jt} \mid s_{jt}, \lambda_t, \theta)$, where a_{jt} summarizes the action of an incumbent. Similarly, the conditional choice probabilities for potential entrants are defined as $\Pr_t\left(a_{jt}^{pe} \mid s_{jt}, \lambda_t, \theta\right)$, with a_{jt}^{pe} summarizing the action of a potential entrant. The likelihood function is then computed as follows:

$$\mathcal{L}(a;\theta) = \prod_{t=1}^{T}\left[\prod_{j \in J_t} \Pr_t\left(a_{jt} | s_{jt}, \lambda_t, \theta\right) \prod_{j \in J_t^{pe}} \Pr_t\left(a_{jt}^{pe} | s_{jt}, \lambda_t, \theta\right)\right], \tag{22}$$

where a summarizes the action profile of all producers observed over the sample period.

The maximum likelihood estimates are obtained by solving

$$\widehat{\theta} = \arg\max_{\theta}\ \ln\left[\mathcal{L}(a;\theta)\right]. \tag{23}$$

Table 2 reports the recovered parameters from the maximum likelihood estimation during the Great Recession. Because this estimation incorporates price and quantity information disciplined by the data, the estimated dynamic costs (values) are expressed in dollar units.[20] The results

[20] Estimates for other industries are reported in Appendix F.

Table 2: Estimates of Dynamic Parameters

Parameters	Notation	Cereal	Yogurt
Mean entry cost ($mn)	γ^{en}	855.16	882.63
		(268.47)	(327.26)
Mean exit value ($mn)	γ^{ex}	85.91	36.58
		(27.71)	(12.96)
Fixed cost of market coverage adjustment ($mn)	γ_0	2.55	1.97
		(0.06)	(0.16)
Linear cost of increasing market coverage ($k)	γ_1^+	55.71	90.31
		(2.70)	(10.65)
Linear cost of decreasing market coverage ($k)	γ_1^-	–9.57	–76.36
		(4.53)	(21.05)
Convex cost of market coverage adjustment	γ_2	1.73	0.91
		(30.66)	(233.75)
Logit shock scale	ψ	1.01	1.15
		(0.09)	(0.05)
Sunk entry cost ($mn)	γ^{en} - γ^{ex}	769.25	846.05

Notes: Standard errors are in parentheses. Results for two industries commonly studied in the literature are presented as illustrative examples. The estimates are based on the sample period 2007Q4–2009Q2, which corresponds to the Great Recession. Results for other sample periods are reported in Appendix F. All dollar values are expressed in 2006Q1 dollars.

indicate that several parameters differ noticeably between the two industries. In particular, exit values are significantly higher in the cereal industry ($85.91 million) than in the yogurt industry ($36.58 million), whereas their entry costs are similar. The sunk cost of entry can be calculated as the mean entry cost minus the mean exit value. Adjustment costs also exhibit heterogeneity across industries, with a notably asymmetric pattern in the yogurt industry: the linear cost of increasing market coverage is approximately $90,000, whereas the value associated with decreasing market coverage (a negative cost) is about $76,000. These large adjustment-cost parameters indicate that such decisions are relatively infrequent in the data.

To assess how well the estimated dynamic parameters explain the observed product market behavior, I evaluate the model's fit by comparing key equilibrium outcomes implied by the model with those observed in the data. Table 3 reports summary statistics of dynamic behaviors, including the average share of new entrants among incumbents, the mean exit rate among incumbents, and the average change in market coverage. The simulated moments closely match their empirical counterparts, indicating that the model successfully replicates producers' dynamic decisions in both

Table 3: Model Fit from Micro Behavior

	Data	Model
Panel A. Cereal		
Entrant share (%)	3.93	3.90
		(0.00)
Exit rate (%)	3.59	3.77
		(0.01)
Changes in market coverage	3.87	3.72
		(0.01)
Panel B. Yogurt		
Entrant share (%)	4.38	4.36
		(0.01)
Exit rate (%)	3.31	3.46
		(0.01)
Changes in market coverage	2.83	2.79
		(0.01)

Notes: Standard errors for the model are calculated using 500 simulations. The model fit is based on the full dynamic estimation sample period from 2007Q4 to 2016Q4.

industries. The close correspondence between data and simulations suggests that the estimated parameters governing entry, exit, and adjustment costs provide an internally consistent explanation of producers' observed dynamic behavior at the product level.

Figure 4 compares the model-simulated aggregate number of products with the observed data for the two industries over the full nonstationary estimation sample (2007Q4–2016Q4). The solid black line represents the data, while the dashed blue line shows the model's simulated path, with the shaded area denoting the 95% confidence intervals based on 500 simulations. The model successfully replicates both the level and the dynamics of product market activity. In the cereal industry (Panel A), the model captures the sharp contraction during the Great Recession and the subsequent persistent upward trend in product variety. In contrast, the yogurt industry (Panel B) exhibits a smoother decline over the same period, which the model also closely tracks.

Figure 4: Model Fit from Aggregate Behavior

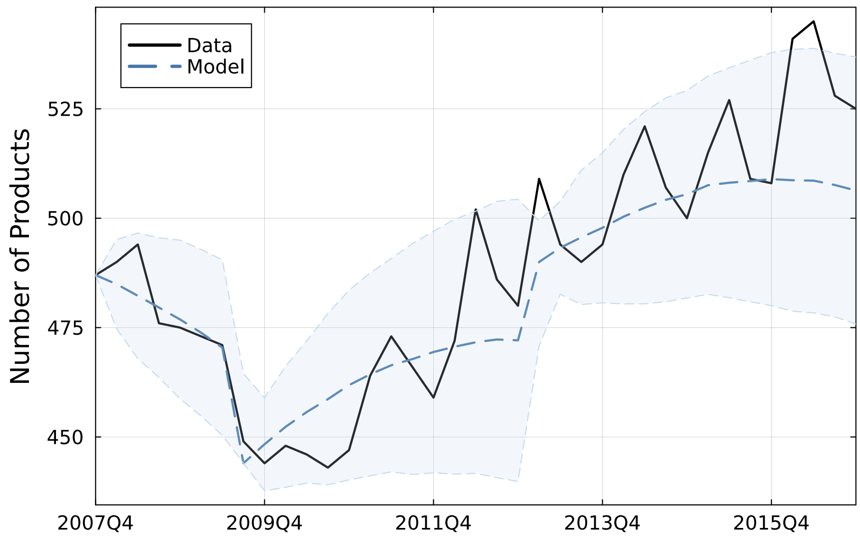

(a) Cereal

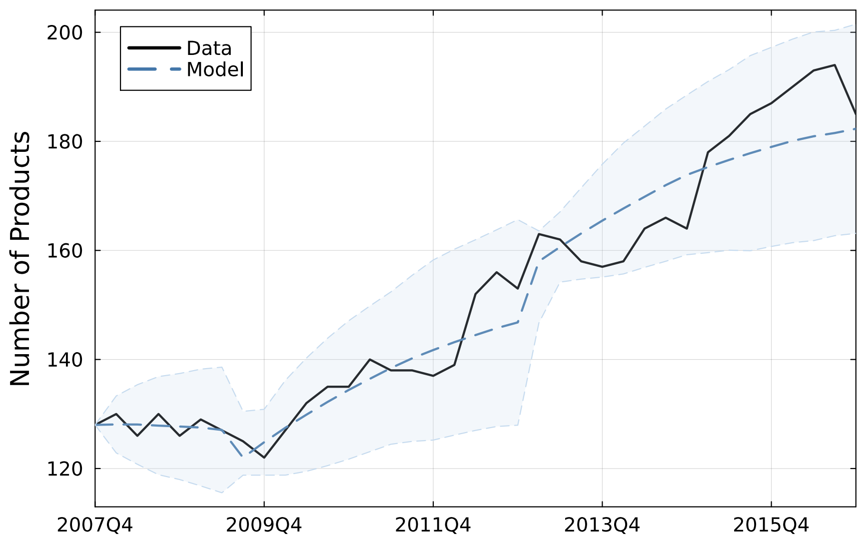

(b) Yogurt

Note: Confidence intervals for the model are calculated using 500 simulations. The model fit is evaluated over the full nonstationary estimation sample period, 2007Q4–2016Q4. For each phase, the number of products in the initial period is directly taken from the data. 95% confidence intervals are shown for the simulated series, with the first-period intervals recalculated using a moving average.

5 Implications of Estimated Model

This section presents empirical patterns on how business cycles affect product market dynamics, emphasizing their connection to demand heterogeneity implied by the estimated model. I relate observed entry, exit, and adjustment behaviors to model-based measures of product appeal and price elasticity of demand. The first subsection examines differences across industries (Section 5.1), while the second focuses on within-industry heterogeneity (Section 5.2).

5.1 Demand Heterogeneity across Industries

Figure 5 illustrates how demand heterogeneity, measured by the median price elasticity of demand in each industry, shapes entry responses over the business cycle. Each dot represents an industry, where the x-axis reports the median price elasticity of demand (derived from the model estimates) and the y-axis shows the percentage change in the number of entrants during the recession relative to the pre-recession period. Industries with less elastic demand (toward the right, with smaller absolute elasticities) experienced a larger decline in entry—up to a 40% reduction—as indicated by the downward-sloping relationship. The intuition is straightforward: when consumers are less sensitive to price changes, entrants face greater difficulty attracting demand even when offering lower prices, as consumers are less willing to substitute toward new products. This challenge is amplified during downturns, when overall market size contracts, leading to a sharper decline in entry.[21]

This relationship can also be explained by the estimated dynamic costs. Figure 6 shows that industries with less elastic demand tend to have higher sunk entry costs: almost twice as large as those in more elastic industries. I calculate the industry-level sunk entry cost as the difference between the mean entry cost and the mean exit value, i.e., $\gamma^{en} - \gamma^{ex}$. Each dot represents an industry, where the x-axis plots the median price elasticity of demand and the y-axis shows the corresponding sunk entry cost. The positive relationship indicates that when consumers are less sensitive to price changes, producers are willing to incur larger upfront costs to enter the market. However, these higher sunk costs, in turn, make entry more difficult and help explain why industries with less elastic demand experience sharper declines in entry during recessions.

[21] Because data on potential entrants are not available and ex ante heterogeneity cannot be observed, the entry analysis is limited to comparisons over time and across industries.

Figure 5: Changes in Product Entry across Industries

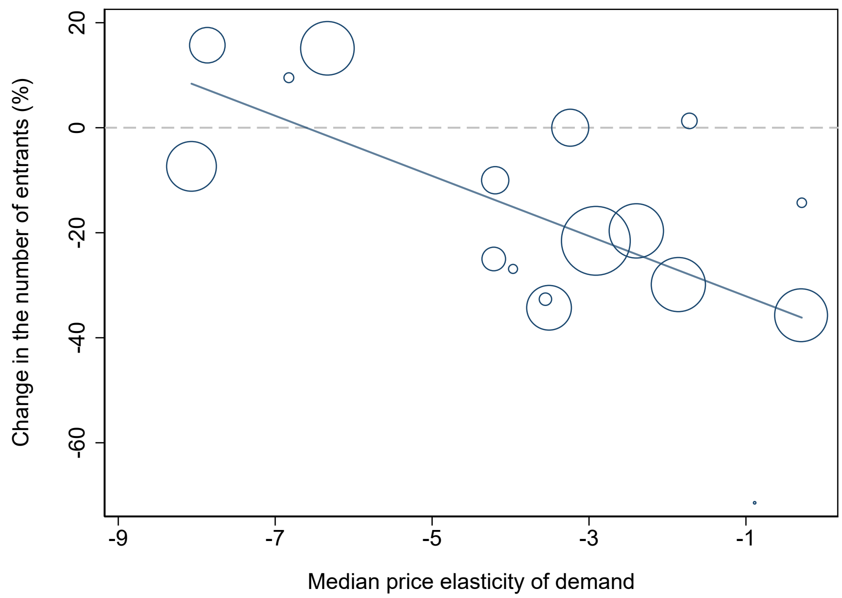

Note: Each dot represents the percentage change in the number of entrants across industries during the Great Recession compared to pre-recession. Dot size reflects total industry revenue.

Lastly, Figure 7 plots changes in the number of exits during the Great Recession relative to the pre-recession period. Unlike entry, the relationship between exit and demand heterogeneity is not as strong. While industries with less elastic demand tend to experience slightly larger increases in exits, the overall pattern remains weak. Nonetheless, some industries exhibit substantially higher exit rates than others. On average, exits increased by about 12% across industries, suggesting that additional factors, such as product-specific characteristics (e.g., product appeal), beyond median price elasticity help explain exit behavior. This contrasts with entry decisions, which are made by ex-ante homogeneous potential entrants. In the next subsection, I examine within-industry heterogeneity in incumbents' dynamic behavior in greater detail.

5.2 Product Heterogeneity within Industry

Figure 8 summarizes how incumbents adjust their dynamic behavior during recessions. Panel (A) shows that exit rates rise sharply by about 0.7 percentage points for products with below-median appeal, indicating that niche products are more vulnerable during downturns. In contrast, high-appeal products remain relatively stable throughout the period. Panel (B) again illustrates this

Figure 6: Sunk Entry Costs across Industries

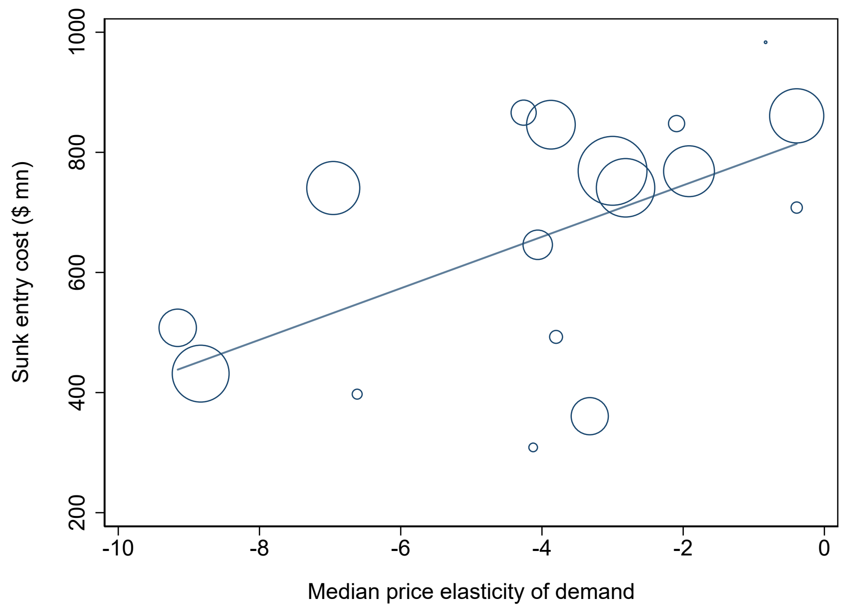

Note: Each dot represents the estimated sunk entry cost for an industry during the Great Recession. Dot size reflects total industry revenue.

asymmetric behavior by showing changes in market coverage, measured by the number of stores (at the chain and region level) carrying each product. Low-appeal products generally reduce their coverage by about three stores, whereas high-appeal products expand it by around two on average. The recession amplifies this divergence: low-appeal products experience a substantial contraction of roughly five stores, while high-appeal products largely sustain their presence. Together, these behaviors imply that recessions intensify selection among incumbents, disproportionately affecting weaker products.

Figure 9 examines how another aspect of demand structure further shapes producers' dynamic decisions during recessions. The figure divides products into elastic and less elastic groups within industries and compares exit rates (Panel A) and changes in market coverage (Panel B) across levels of product appeal. Low-appeal products are more likely to exit, and this effect is particularly pronounced among less elastic products by about one percentage point relative to low-appeal products with elastic demand. The contrast is even stronger in market coverage: low-appeal products with less elastic demand reduce their coverage by an average of seven stores, whereas products with more elastic demand or higher appeal show little change. The intuition behind this pattern is that when low-appeal producers face less-elastic demand, cutting prices does little to

Figure 7: Changes in Product Exit across Industries

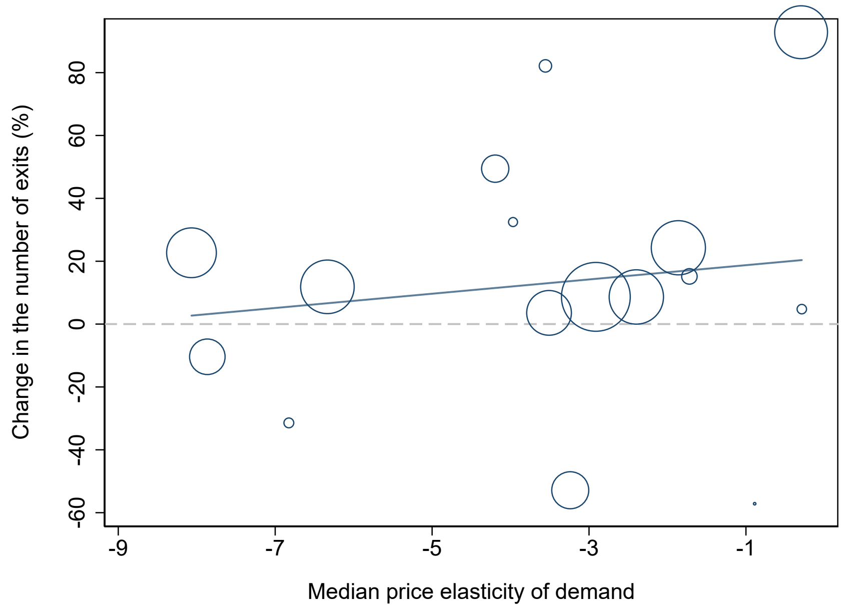

Note: Each dot represents the percentage change in the number of exits across industries during the Great Recession compared to pre-recession. Dot size reflects total industry revenue.

attract new consumers, as they are less responsive to price changes. Consequently, these products are more likely to exit or substantially reduce their market presence during recessions. These results suggest that demand elasticity and product appeal jointly determine producers' dynamic behavior during recessions.

6 Counterfactual Analysis

So far, the analysis has shown that product market dynamics vary substantially across industries, with demand heterogeneity serving as a key driver of these differences. This section conducts counterfactual experiments to quantify the welfare implications of these heterogeneous dynamics. I first examine how business cycle fluctuations affect welfare by measuring the welfare costs of the business cycle (Section 6.1). Next, I evaluate how these welfare costs differ across industries (Section 6.2). Finally, I conduct a counterfactual experiment that removes demand heterogeneity (Section 6.3).

Figure 8: Changes in Exit and Market Coverage Adjustment

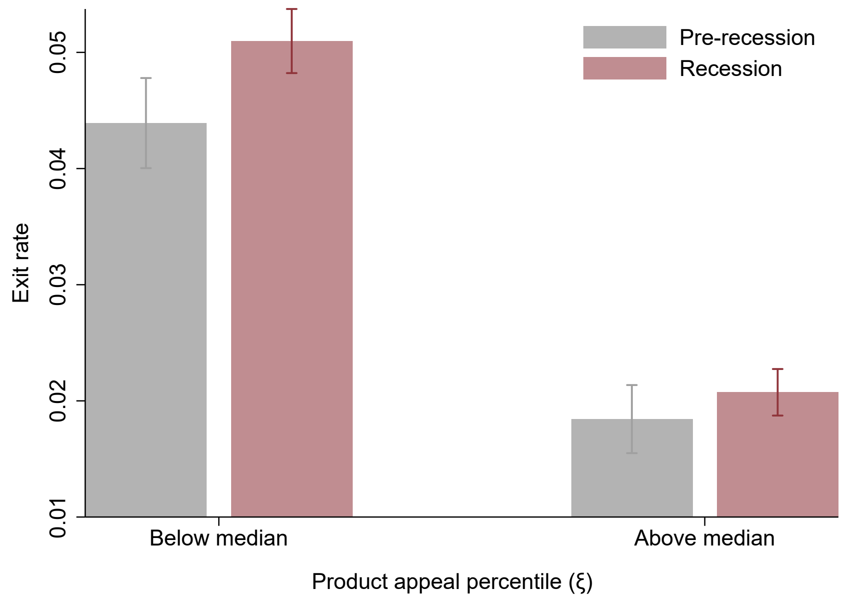

(a) Exit

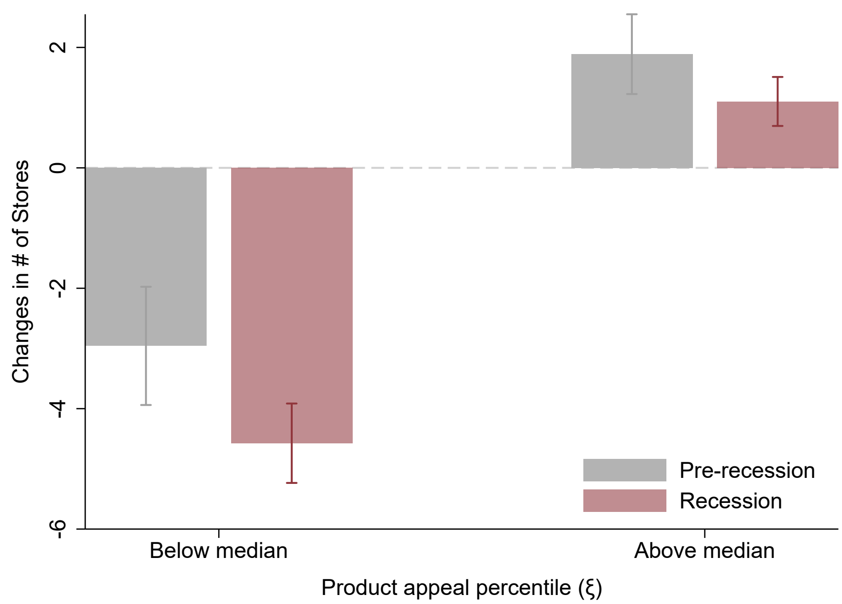

(b) Market Coverage Adjusment

Note: Each bar represents the mean exit rate (Panel A) and the average change in the number of stores (at the chain and region level) carrying each product (Panel B) before and during the Great Recession. Products are divided into below- and above-median groups based on recovered product appeal (ξ). The pre-recession period covers 2007Q1–2007Q3, and the recession period covers 2007Q4–2009Q2. Error bars indicate 90% confidence intervals.

Figure 9: Mean Exit and Market Coverage Adjustment

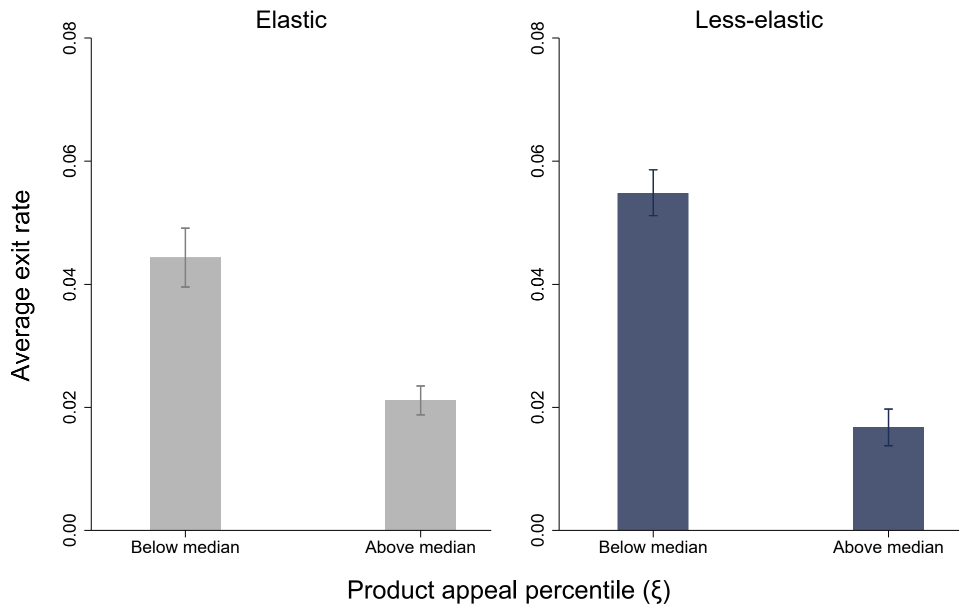

(a) Exit

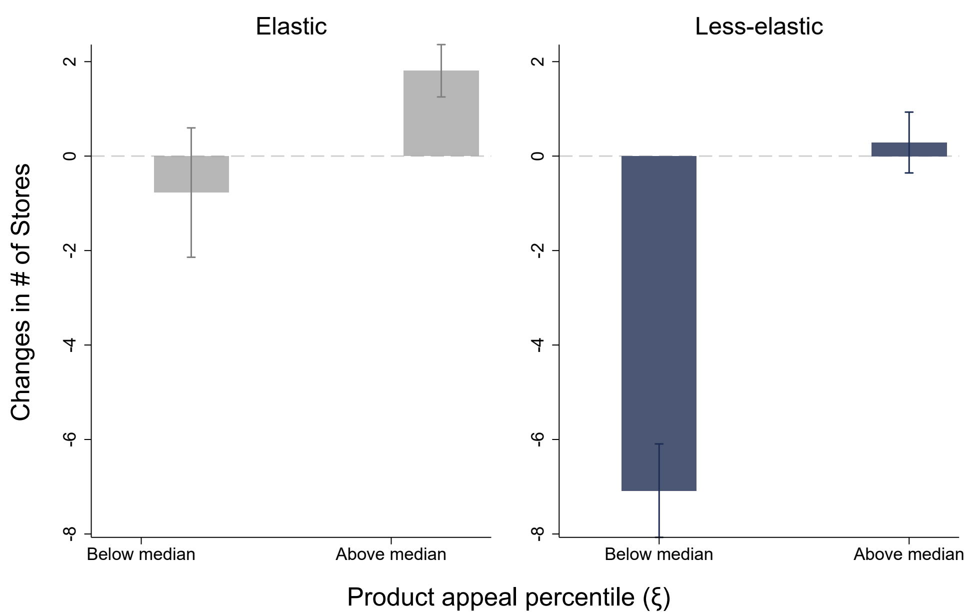

(b) Market Coverage Adjusment

Note: Each bar represents the mean exit rate (Panel A) and the average change in the number of stores carrying each product (Panel B) during the Great Recession. Products are divided into below- and above-median groups based on estimated product appeal (ξ) and into elastic versus less-elastic groups according to the recovered product-level price elasticity. The recession period covers 2007Q4–2009Q2. Error bars indicate 90% confidence intervals.

6.1 Welfare Costs of the Business Cycle

First, I compute consumer and producer surplus at time t using the estimated parameters from the previous section. Consumer surplus, CS_t, is defined as

$$CS_t = \int \left\{ \frac{\gamma_E + \ln\left[1 + \sum_{j \in J_t} \exp\left(\alpha_i p_{jt} + \phi \log N_{jt} + \xi_{jt}\right)\right]}{-\alpha_i} \right\} dF(\alpha_i) \cdot M_t \cdot \exp(z_t), \tag{24}$$

where γ_E denotes the Euler constant (≈ 0.57721). The term divided by $-\alpha_i$ converts utility units into monetary value, consistent with the inclusive value formulation.

Similarly, producer surplus, PS_t, is defined as

$$\begin{aligned} PS_t = &\sum_{j \in J_t} \left\{ \pi_t(s_{jt}; \alpha, \sigma_\alpha^2, \phi) + \mathbb{1}_{\{a_{jt}=0\}} \gamma_{jt}^{ex} - \mathbb{1}_{\{a_{jt}>0\}} \left[C(\Delta_N; \gamma^c) - \nu_{jt} \right] \right\} \\ &- \sum_{j \in J_t^{pe}} \mathbb{1}_{\{a_{jt}^{pe}=1\}} \gamma_{jt}^{en}, \end{aligned} \tag{25}$$

where PS_t represents the expected profit of all producers in the market, net of entry, exit, and adjustment costs.

Based on equations (24) and (25), I compute social welfare, $\mathbf{SW}$, as the sum of consumer and producer surplus over the entire sample period, pooling across all industries in the sample and applying a social discount factor β_s. The social discount factor is set higher than the private discount factor β to reflect the social planner's longer-term perspective. Specifically, I use $\beta_s = 0.999$ and $\beta = 0.99$. The social welfare is then expressed as

$$\mathbf{SW} = \mathbb{E}\left[\sum_{t \geq 1}^{\bar{T}} \beta_s^{t-1} (CS_t + PS_t) \right]. \tag{26}$$

I conduct an exercise comparing welfare outcomes under two scenarios: the baseline economy with aggregate demand shocks and a counterfactual economy without shocks. For each scenario, I simulate 500 independent paths of industry evolution and compute social welfare, consumer surplus, and producer surplus as described above. Table 4 reports the average welfare across these simulations. The welfare cost of the business cycle is measured as the difference between the two scenarios, and the corresponding standard errors are reported in parentheses. The results

Table 4: Welfare Costs of the Business Cycle

	Social Welfare	Consumer Surplus	Producer Surplus
With shock	2,873.27	260.42	2,612.85
Without shock	2,910.78	301.84	2,608.94
Welfare costs (difference)	-37.51	-41.42	3.90
	(0.49)	(0.07)	(0.48)

Notes: The table reports welfare measures under two scenarios: the baseline economy with aggregate demand shocks and the counterfactual economy without shocks. The last row shows the difference (with shock minus without shock). All values are in billions of dollars. Standard errors are reported in parentheses.

indicate that aggregate fluctuations generate a substantial welfare loss of about $37 billion, which corresponds to approximately 1.5% of pre-recession personal consumption expenditure (PCE) on food. Given that the NielsenIQ data represent only a subset of the Consumer Price Index (CPI) basket and that the current sample includes roughly 15 industries, this estimated welfare cost is considerably large.

The welfare loss arises primarily through reduced consumer surplus, which declines by about $41 billion, while producer surplus remains relatively stable. Given that consumer surplus accounts for roughly 10% of total social welfare, this represents a sizable impact. The underlying mechanism is that recessions lead to fewer product entries and more exits, resulting in a market composed of fewer but higher-appeal products offered at higher prices. The decline in consumer surplus indicates that the gains from access to higher-appeal products do not fully offset the losses from higher prices and reduced product variety.

6.2 Welfare Costs across Industries

Next, I turn to the exercise at the disaggregated industry level. Figure 10 illustrates the welfare costs across industries with different demand elasticities. Panel A presents the change in consumer surplus (ΔCS), while Panel B shows the corresponding change in producer surplus (ΔPS). The x-axis in both panels represents the median price elasticity of demand for each industry, derived from the estimated model. The results show that industries with less elastic demand experience larger declines in consumer surplus during recessions. This pattern arises because substitution possibilities are limited in these industries, so consumers continue purchasing similar products even as prices rise, leading to greater welfare losses. For producer surplus, industries with less elastic demand tend to experience slight increases following shocks. This pattern is mainly driven by higher exit rates in those industries, as producers obtain larger exit values and face lower effective entry costs.

6.3 Counterfactual Exercise: No Consumer Heterogeneity

Finally, I examine the role of demand-side heterogeneity in shaping welfare outcomes through product market dynamics. In the baseline model, consumers differ in their price sensitivity within each industry. To assess the importance of this heterogeneity, I conduct a counterfactual exercise setting the variance of price sensitivity to zero ($\sigma_\alpha^2 = 0$), which collapses the random-coefficients logit model into a standard logit demand system, closely resembling the representative-consumer framework in a CES demand model (Anderson, De Palma, and Thisse, 1992). I still allow for differences across industries in the mean price elasticity (α), consistent with a nested CES structure that captures across-industry heterogeneity. This experiment isolates how ignoring within-industry consumer heterogeneity affects the magnitude of the welfare costs of the business cycle.

The results in Table 5 show that the welfare costs of the business cycle are understated by roughly 5% when consumer heterogeneity is ignored, equivalent to about 0.2% of pre-recession personal consumption expenditure (PCE) on food. The attenuation arises because a representative consumer model overestimates substitutability across products, thereby dampening the welfare impact of product entry and exit. In contrast, when consumers differ in their preferences, changes in product variety and composition have stronger welfare implications, as some consumers are more sensitive to changes in the set of available products through entry and exit. These findings underscore the importance of accounting for heterogeneous consumer responses when evaluating welfare dynamics in differentiated-product markets.

7 Conclusion

This paper investigates how business cycles reshape product market dynamics, focusing on the interaction between consumer demand and producer entry and exit. I develop a dynamic model with heterogeneous producers and consumers and estimate it using detailed micro-level scanner

Figure 10: Changes in Welfare across Industries

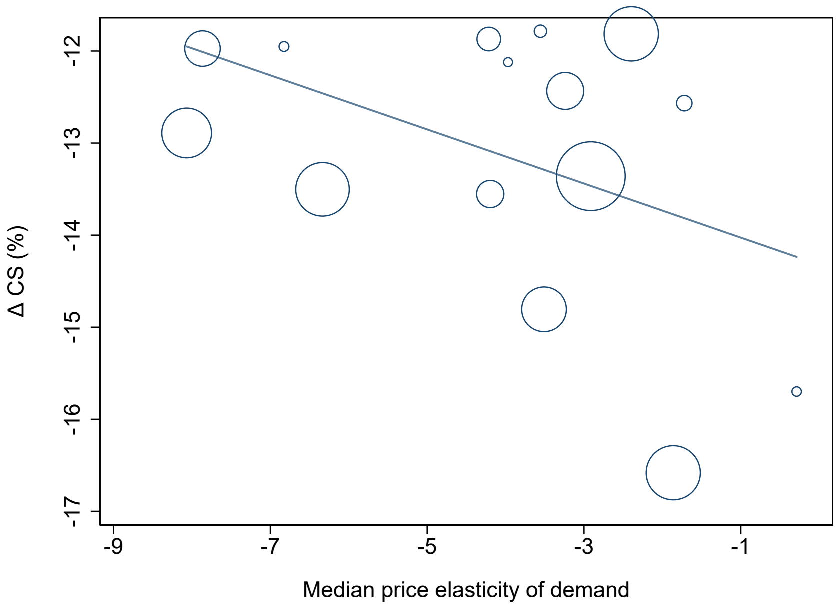

(a) Consumer Surplus

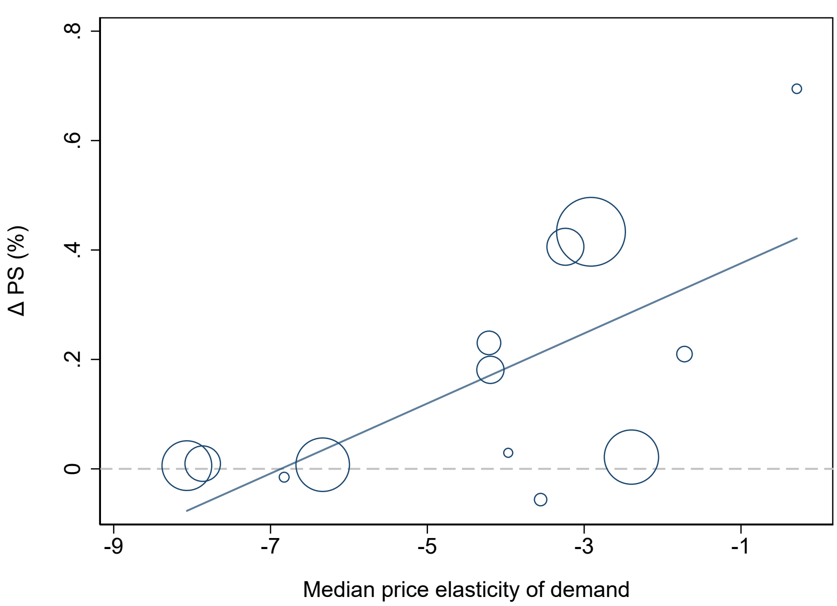

(b) Producer Surplus

Notes: Each bubble represents an industry, with bubble size proportional to its revenue. The x-axis shows the median price elasticity of demand derived from the estimated model. ΔCS and ΔPS denote percentage changes in consumer surplus and producer surplus, respectively, between the baseline (with shock) and the counterfactual (without shock) economies.

Table 5: Welfare Costs under No Consumer Heterogeneity

	ΔSW	ΔCS	ΔPS
Baseline	-37.51	-41.42	3.90
	(0.49)	(0.07)	(0.48)
No consumer heterogeneity	-35.59	-39.23	3.64
	(0.45)	(0.07)	(0.44)
Compared to baseline (%)	-5.12	-5.27	-6.77

Notes: The table reports welfare costs for each scenario. The last row shows the percentage difference of the counterfactual relative to the baseline. All values are in billions of dollars. Standard errors are reported in parentheses.

data from the U.S. consumer goods sector. The estimated model provides new evidence on how aggregate demand shocks affect product entry and exit decisions, highlighting the critical role of demand heterogeneity in shaping these dynamics and their welfare implications.

The analysis shows that producers systematically reduce product entry and increase exit during recessions, with the magnitude of these adjustments varying across markets. These differences arise from variation in the price elasticity of demand and product characteristics. For entry, markets with less elastic demand, characterized by higher sunk entry costs, experience particularly large declines in product introduction. For exit behavior, products with lower consumer appeal and facing less elastic demand are more likely to exit and reduce their market coverage during recessions. This finding is important because, in a standard monopolistic competition framework, the price elasticity of demand is often viewed as the inverse of the markup, suggesting greater market power in less elastic markets. However, this paper shows that such an interpretation can be misleading: less elastic demand is associated with higher markups only for high-appeal products that remain in the market, whereas low-appeal products facing limited substitution possibilities are more likely to exit despite charging lower prices.

These dynamics imply significant welfare losses for consumers, primarily through higher prices and reduced product variety. The welfare cost of the business cycle is substantial, and this decline in welfare is largely driven by the reduction in consumer surplus. Moreover, the magnitude of welfare losses varies systematically across industries, with those facing less elastic demand experiencing larger declines. Finally, ignoring consumer heterogeneity leads to an underestimation of welfare costs by about 5%, highlighting the importance of accounting for diverse consumer preferences when evaluating welfare dynamics in differentiated-product markets.

The heterogeneity in product market responses underscores the importance of accounting for rich demand structures when analyzing business cycle effects. Future research could extend this framework to examine how consumer demand itself evolves over the business cycle, particularly through trading-down behavior, or to incorporate non-homothetic preferences into the model. Overall, this paper highlights the importance of studying product market dynamics as a critical channel through which business cycles affect welfare.

References

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky. 2019. "Regressive sin taxes, with an application to the optimal soda tax." *The Quarterly Journal of Economics* 134(3):1557–1626.

Alvarez, Fernando, and Urban J Jermann. 2004. "Using asset prices to measure the cost of business cycles." *Journal of Political economy* 112(6):1223–1256.

Anderson, Simon P, Andre De Palma, and Jacques-Francois Thisse. 1992. *Discrete choice theory of product differentiation*. MIT press.

Argente, David, and Munseob Lee. 2021. "Cost of living inequality during the great recession." *Journal of the European Economic Association* 19(2):913–952.

Argente, David, Munseob Lee, and Sara Moreira. 2018. "Innovation and product reallocation in the great recession." *Journal of Monetary Economics* 93:1–20.

Atalay, Enghin, Erika Frost, Alan T Sorensen, Christopher J Sullivan, and Wanjia Zhu. 2023. "Scalable demand and markups."Technical report, National Bureau of Economic Research.

Atkeson, Andrew, and Christopher Phelan. 1994. "Reconsidering the costs of business cycles with incomplete markets." *NBER macroeconomics annual* 9:187–207.

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2021. "Common ownership and competition in the ready-to-eat cereal industry."Technical report, National Bureau of Economic Research.

Becker, Jonathan. 2024. "Do Poor Households Pay Higher Markups in Recessions?"Technical report.

Benkard, C Lanier. 2004. "A dynamic analysis of the market for wide-bodied commercial aircraft." *The Review of Economic Studies* 71(3):581–611.

Benkard, C Lanier, Przemyslaw Jeziorski, and Gabriel Weintraub. 2024. "Transitional Market Dynamics in Complex Environments."Technical report, National Bureau of Economic Research.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*:841–890.

Berry, Steven T. 1994. "Estimating discrete-choice models of product differentiation." *The RAND Journal of Economics*:242–262.

Bilbiie, Florin O, Fabio Ghironi, and Marc J Melitz. 2012. "Endogenous entry, product variety, and business cycles." *Journal of Political Economy* 120(2):304–345.

Bils, Mark, and Peter J Klenow. 1998. "Using consumer theory to test competing business cycles models." *Journal of Political Economy* 106(2):233–261.

Brand, James. 2021. "Differences in differentiation: Rising variety and markups in retail food stores." *Available at SSRN 3712513*.

Broda, Christian, and David E Weinstein. 2010. "Product creation and destruction: Evidence and price implications." *American Economic Review* 100(3):691–723.

Charles, Kerwin, and Melvin Stephens Jr. 2006. "The level and composition of consumption over the business cycle: the role of" quasi-fixed" expenditures."

Chatterjee, Satyajit, and Dean Corbae. 2007. "On the aggregate welfare cost of Great Depression unemployment." *Journal of monetary Economics* 54(6):1529–1544.

Clementi, Gian Luca, and Berardino Palazzo. 2016. "Entry, exit, firm dynamics, and aggregate fluctuations." *American Economic Journal: Macroeconomics* 8(3):1–41.

Collard-Wexler, Allan. 2013. "Demand fluctuations in the ready-mix concrete industry." *Econometrica* 81(3):1003–1037.

Conlon, Christopher, and Jeff Gortmaker. 2023. "Incorporating micro data into differentiated products demand estimation with PyBLP."Technical report, National Bureau of Economic Research.

Constantinides, George M. 2025. "Welfare Costs of Idiosyncratic and Aggregate Consumption Shocks." *The Review of Asset Pricing Studies* 15(2):103–120.

Corbae, Dean, and Pablo D'Erasmo. 2021. "Capital buffers in a quantitative model of banking industry dynamics." *Econometrica* 89(6):2975–3023.

De Santis, Massimiliano. 2007. “Individual consumption risk and the welfare cost of business cycles.” *American Economic Review* 97(4):1488–1506.

DellaVigna, Stefano, and Matthew Gentzkow. 2019. “Uniform pricing in us retail chains.” *The Quarterly Journal of Economics* 134(4):2011–2084.

Dolmas, Jim. 1998. “Risk preferences and the welfare cost of business cycles.” *Review of Economic Dynamics* 1(3):646–676.

Döpper, Hendrik, Alexander MacKay, Nathan H Miller, and Joel Stiebale. 2025. “Rising markups and the role of consumer preferences.” *Journal of Political Economy* 133(8):000–000.

Dubois, Pierre, Rachel Griffith, and Martin O’Connell. 2022. “The use of scanner data for economics research.” *Annual Review of Economics* 14:723–745.

Dunne, Timothy, Shawn D Klimek, Mark J Roberts, and Daniel Yi Xu. 2013. “Entry, exit, and the determinants of market structure.” *The RAND Journal of Economics* 44(3):462–487.

Epaulard, Anne, and Aude Pommeret. 2003. “Recursive utility, endogenous growth, and the welfare cost of volatility.” *Review of Economic Dynamics* 6(3):672–684.

Ericson, Richard, and Ariel Pakes. 1995. “Markov-perfect industry dynamics: A framework for empirical work.” *The Review of Economic Studies* 62(1):53–82.

Etro, Federico, and Andrea Colciago. 2010. “Endogenous market structures and the business cycle.” *The Economic Journal* 120(549):1201–1233.

Gali, Jordi, Mark Gertler, and J David Lopez-Salido. 2007. “Markups, gaps, and the welfare costs of business fluctuations.” *The review of economics and statistics* 89(1):44–59.

Gamber, William. 2023. “The Importance of Entry in the Business Cycle: What Are the Roles of Markups, Adjustment Costs, and Heterogeneity?” *Journal of Political Economy Macroeconomics* 1(3):551–585.

Gandhi, Amit, and Jean-François Houde. 2019. “Measuring substitution patterns in differentiated-products industries.” *NBER Working paper*(w26375).

Granja, Joao, and Sara Moreira. 2023. "Product innovation and credit market disruptions." *The Review of Financial Studies* 36(5):1930–1969.

Hamano, Masashige, and Francesco Zanetti. 2017. "Endogenous product turnover and macroeconomic dynamics." *Review of Economic Dynamics* 26:263–279.

Hausman, Jerry A. 1996. *Valuation of new goods under perfect and imperfect competition*. National Bureau of Economic Research Cambridge, Mass., USA.

Hopenhayn, Hugo A. 1992. "Entry, exit, and firm dynamics in long run equilibrium." *Econometrica* 60(5):1127–1150.

Hottman, Colin J, Stephen J Redding, and David E Weinstein. 2016. "Quantifying the sources of firm heterogeneity." *The Quarterly Journal of Economics* 131(3):1291–1364.

Hristakeva, Sylvia. 2022. "Vertical Contracts with Endogenous Product Selection: An Empirical Analysis of Vendor Allowance Contracts." *Journal of Political Economy* 130(12):3202–3252.

Ifrach, Bar, and Gabriel Y Weintraub. 2017. "A framework for dynamic oligopoly in concentrated industries." *The Review of Economic Studies* 84(3):1106–1150.

Igami, Mitsuru. 2017. "Estimating the innovator's dilemma: Structural analysis of creative destruction in the hard disk drive industry, 1981–1998." *Journal of Political Economy* 125(3): 798–847.

Imrohoroğlu, Ayşe. 1989. "Cost of business cycles with indivisibilities and liquidity constraints." *Journal of Political economy* 97(6):1364–1383.

Jaimovich, Nir, and Max Floetotto. 2008. "Firm dynamics, markup variations, and the business cycle." *Journal of Monetary Economics* 55(7):1238–1252.

Jaimovich, Nir, Sergio Rebelo, and Arlene Wong. 2019. "Trading down and the business cycle." *Journal of Monetary Economics* 102:96–121.

Jeon, Jihye. 2022. "Learning and investment under demand uncertainty in container shipping." *The RAND Journal of Economics* 53(1):226–259.

Krebs, Tom. 2003. “Growth and welfare effects of business cycles in economies with idiosyncratic human capital risk.” *Review of Economic Dynamics* 6(4):846–868.

——— 2007. “Job displacement risk and the cost of business cycles.” *American Economic Review* 97(3):664–686.

Krusell, Per, and Anthony A Smith, Jr. 1998. “Income and wealth heterogeneity in the macroeconomy.” *Journal of political Economy* 106(5):867–896.

Krusell, Per, and Anthony A Smith Jr. 1999. “On the welfare effects of eliminating business cycles.” *Review of Economic Dynamics* 2(1):245–272.

Lee, Yoonsoo, and Toshihiko Mukoyama. 2015. “Entry and exit of manufacturing plants over the business cycle.” *European Economic Review* 77:20–27.

Lucas, Robert E. Jr. 1987. *Models of Business Cycles*. New York Basil Blackwell.

Michelacci, Claudio, Luigi Paciello, and Andrea Pozzi. 2022. “The extensive margin of aggregate consumption demand.” *The Review of Economic Studies* 89(2):909–947.

Nevo, Aviv. 2001. “Measuring market power in the ready-to-eat cereal industry.” *Econometrica* 69(2):307–342.

——— 2011. “Empirical models of consumer behavior.” *Annu. Rev. Econ.* 3(1):51–75.

Obstfeld, Maurice. 1994. “Evaluating risky consumption paths: The role of intertemporal substitutability.” *European economic review* 38(7):1471–1486.

Orchard, Jacob. 2022. “Cyclical demand shifts and cost of living inequality.” *Available at SSRN 4033572*.

Otrok, Christopher. 2001. “On measuring the welfare cost of business cycles.” *Journal of Monetary Economics* 47(1):61–92.

Reis, Ricardo. 2009. “The Time-Series Properties of Aggregate Consumption: Implications for the Costs of Fluctuations.” *Journal of the European Economic Association* 7(4):722–753.

Rust, John. 1987. "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher." *Econometrica: Journal of the Econometric Society*:999–1033.

Salanié, Bernard, and Frank A Wolak. 2022. "Fast, detail-free, and approximately correct: Estimating mixed demand systems."Technical report, Working paper.

Seim, Katja. 2006. "An empirical model of firm entry with endogenous product-type choices." *The RAND Journal of Economics* 37(3):619–640.

Siemer, Michael. 2014. "Firm entry and employment dynamics in the great recession." FEDS Working Paper no. 2014-56, Federal Reserve Board. Washington, DC.

Storesletten, Kjetil, Chris I Telmer, and Amir Yaron. 2001. "The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk." *European Economic Review* 45(7):1311–1339.

Thurk, Jeff. 2024. "Outsourcing, firm innovation, and industry dynamics in the production of semiconductors." *Working Paper*.

Wollmann, Thomas G. 2018. "Trucks without bailouts: Equilibrium product characteristics for commercial vehicles." *American Economic Review* 108(6):1364–1406.

Appendix

A Details on the Data

This section provides additional details on the construction of the dataset used in the analysis. Section A.1 describes how the scanner data are processed to form the analysis sample, including the aggregation procedures and product classification. Section A.2 documents the supplementary macro-level data sources, and Section A.3 outlines the sample selection criteria applied in the empirical analysis.

A.1 Data Construction

Product definition This paper relies on the NielsenIQ Retail Scanner Data, which provide weekly sales information at the UPC (Universal Product Code) level for individual retail chain stores. Each UPC corresponds to a unique product variant—for example, a 2-liter bottle of *Diet Pepsi*. In this paper, I define a product at the brand level, aggregating all UPCs belonging to the same brand. Under this definition, the product *Diet Pepsi* encompasses multiple UPCs that represent different package sizes and formats, such as a 2-liter bottle or an eight-pack of 12-ounce cans. The brand classification in this dataset is narrowly defined. For instance, *Coca-Cola Cherry* and *Coca-Cola Zero Sugar* are treated as distinct products rather than variants of *Coca-Cola*. The number of such branded products per industry in the sample averages about 280, with a median of 160.

To conduct this aggregation, I measure quantities as the total number of units sold, converted into a standard unit of measurement (such as ounces) following the official data documentation. Prices are defined on the same unit basis (e.g., dollars per ounce). For each industry, I adopt the most common unit of measurement as the standard. In the industries included in my current sample, the unit of measure is generally unique within each industry.

Figure A1 illustrates the distribution of the number of UPCs within each brand, pooled across industries in the sample. The figure shows that a substantial share of products consist of a single UPC, corresponding to one specific size or multipack configuration. However, a subset of products includes multiple UPCs that differ in size or packaging format. The median number of UPCs per brand is 2, and the mean is 5. Overall, brand-level aggregation remains sufficiently granular, as most differentiation occurs across brands rather than across UPC variants within a brand. In this sense, a brand can be viewed as a close approximation to a product defined at the UPC level.

Figure A1: Number of UPCs within Brand

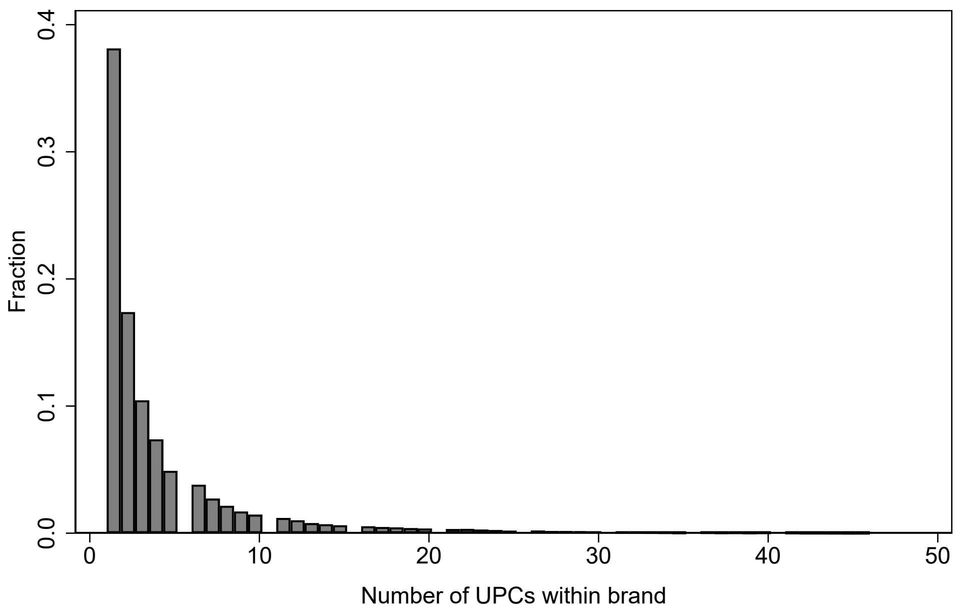

Note: The figure shows the distribution of the number of UPCs (Universal Product Code) within each brand, pooled across industries in the sample.

Industry definition NielsenIQ provides detailed product categories, each of which can be considered a well-defined product market. Accordingly, this paper defines an industry as one of these data-driven product markets. Based on this definition, the analysis is performed separately for each industry. For this reason, the current analysis does not consider cross-category substitution. The industries included in the current analysis are: canned soup, pasta (spaghetti), ready-to-eat cereals, cookies, jams, honey, peanut butter, bottled water, grated cheese, refrigerated yogurt, specialty soap, hand cleaners and sanitizers, tooth cleaners, deodorant, disposable diapers, and pre-shave cosmetics. Disposable diapers and pre-shave cosmetics are excluded from the dynamic estimation due to limited observations in the data.

While the selection is partly informed by Döpper et al. (2025), who provide a revenue-based accounting of product categories in the dataset, it is not identical to theirs. I use their information as a reference and additionally consider factors such as product durability and purchase frequency when determining which industries to include. In Section F, I show that the selected industries exhibit substantial variation in the median elasticity of demand. These industries provide sufficient heterogeneity in demand characteristics, which are central to the analysis in this paper.

Local market definition NielsenIQ Retail Scanner Data are reported at the UPC$\times$week$\times$store level. I aggregate these observations to the retail chain$\times$region$\times$quarter level. A chain corresponds to the parent retail company, and a region is defined at the DMA (Designated Market Area) level. Throughout the paper, references to a *local market* or *store* refer to this aggregated chain$\times$region market, unless otherwise specified as the actual store level. This local market serves as the unit of analysis in the demand estimation. Variation across these local markets is also used to measure market coverage—that is, the number of stores in which each product is sold. Prices for each brand (aggregated across UPCs) are computed as quantity-weighted averages across individual stores within each local market, and quantities are defined as total sales volume (in standardized measurement units) summed across stores and UPCs within the same market.

Market definition In most cases, this paper uses the terms *product market* (or simply, *market*) and *industry* interchangeably. This is consistent with the industry definition provided above. This level serves as the unit of the dynamic estimation, and entry and exit are defined at this level on a quarterly basis. After conducting the demand estimation at the local-market level, I aggregate key variables to the national level for each industry over time, applying quantity weights in this process. This aggregation allows me to recover the state variable, *product appeal*, from the demand estimation and incorporate it into the dynamic model. Prices, which are used to compute profits in the dynamic estimation, are also aggregated to the national level using the same procedure. Market coverage refers to the number of local markets in which each product is sold. The dynamic analysis sample therefore consists of a brand-by-quarter panel, where each observation includes the national-level price and quantity sold, the number of local markets where the product is available, indicators for entry and exit decisions, and the product appeal recovered from the demand estimation.

Market size In this paper, I assume that there are M_t consumers in each market, which is required for defining market shares in the demand estimation. The literature typically uses population size as a proxy for market size, and I closely follow the approach in Döpper et al. (2025). Specifically, I calibrate market size so that the combined share of inside goods centers around 0.45 (with a maximum of 0.9), consistent with their procedure. To construct M_t, I use annual county-level population data and multiply it by the number of actual stores within each retail chain and DMA region for each quarter. This is the only step where I use disaggregated store counts from the raw data (i.e., the actual number of observed stores). Hence, market size in the demand estimation refers to this measure at the local-market level.

For the dynamic estimation, market size remains relevant because the model continues to account for market shares over time at the national level. In this case, I aggregate the local-market populations and store counts across all regions and chains to construct the national market size for each period. I then apply the same calibration procedure as in Döpper et al. (2025), ensuring that the sum of total product market shares observed in the national brand-level panel is centered around 0.45.

In this paper, I assume that aggregate shocks are the only factors that affect market size over time in order to capture business cycle effects. However, the data show that national market size increases over time—for example, due to large migration flows into certain counties. To isolate shock-driven variation in the dynamic estimation, I calibrate the market size in the dynamic model to the pre-recession average level and hold it constant over time. Consequently, only aggregate demand shocks influence market size, allowing the model to capture cyclical fluctuations in overall demand without being confounded by long-run structural changes.

Comparison to consumer panel NielsenIQ also provides the Consumer Panel dataset, which has been widely used in related research. The Consumer Panel records purchase information from a sample of households who self-report their transactions, resulting in limited product coverage and smaller sample size. In contrast, the Retail Scanner dataset captures billions of transactions directly at the point of sale, providing comprehensive coverage across retail chains and product categories. While the scale of the Retail Scanner data imposes higher computational requirements, it enables a more detailed characterization of product-market dynamics and market structure over time.

A.2 Other Data Sources and Construction Details

To construct real prices, I incorporate macroeconomic series from the Federal Reserve Economic Data (FRED). Specifically, I use the quarterly *Consumer Price Index for All Urban Consumers: All Items Less Food and Energy in U.S. City Average* (CPILFESL) to deflate nominal prices to real terms, expressed in 2006Q1 dollars.

For the construction of aggregate demand shocks, I use the national unemployment rate (UNRATE) from FRED, converting the monthly series into quarterly averages. I then define the aggregate shock as the negative deviation of the unemployment rate from its pre-recession mean, so that higher unemployment corresponds to a more negative demand shock. This normalized series is directly used in the dynamic estimation as an exogenous state variable. Because the nonstationary model assumes convergence toward a stationary equilibrium after sufficiently long periods, I extend the analysis by 30 additional quarters beyond the final observed period. These extended periods are not used in the estimation but are included only for solving the model. For these appended periods, the aggregate shock is held constant at the level observed in the final sample period.

To prevent unrealistically large fluctuations in aggregate demand when solving the model, I apply a sensitivity adjustment to the aggregate shock. The adjustment is parameterized as $\exp(\gamma z_t)$, where z_t denotes the normalized shock and γ is a sensitivity parameter set to 0.05. The value of γ is calibrated from a reduced-form regression of quantities on the constructed demand shock, ensuring that the magnitude of the model-implied responses remains consistent with those in the data.

For the instrumental-variable construction, I additionally use the local unemployment rate as an exogenous instrument (Backus, Conlon, and Sinkinson, 2021). This variable, drawn from the BLS *Local Area Unemployment Statistics* (LAUS), captures county-level demand conditions that vary across regions and over time. Incorporating this measure provides plausibly exogenous variation that strengthens the identification of demand parameters across markets.

A.3 Sample Restriction

Several restrictions are applied to construct a sample. First, the analysis is limited to grocery stores, drug stores, and mass merchandisers, which together provide the most representative coverage of nationwide sales in the dataset. Other outlet types, such as convenience stores and liquor stores, are excluded because they are not comprehensively represented in the data.

Second, only stores that are continuously observed from 2006 to 2019 within each retail chain and region are included. This restriction avoids confounding effects from store entry and exit over time and ensures that producers' market coverage decisions are made within a fixed set of local markets. To construct the balanced-store sample, I begin with the raw store-level data and identify stores that are present in all periods between 2006 and 2019.

Third, to avoid spurious classification of entry and exit, brief gaps in observed sales are not treated as market exit. Specifically, if a product is missing for up to three consecutive quarters, it is classified as a continuing incumbent (i.e., no exit), although those missing quarters are excluded from the estimation. If a product reappears after four or more consecutive missing quarters, it is classified as a new entry. This adjustment affects only a very small share of cases.

Lastly, private-label products are included in the demand estimation but excluded from the dynamic estimation. The market share of these private-label brands (retailer-specific products) is substantial in the data, making it important to account for them when estimating demand. After estimating the demand parameters with them, I treat the presence of private labels as exogenous in the dynamic estimation and therefore exclude them. This exclusion is also motivated by data limitations: retailer identifiers are restricted, making it impossible to observe or model private-label entry and exit decisions.

This paper does not impose additional sample restrictions. Accounting for product entry and exit requires capturing the full set of products available in the data. For example, Döpper et al. (2025) focus on the top 20 brands and treat the remaining brands as fringe products. In contrast, I include all observed products, as entry and exit occur primarily among these niche brands. As a result, their estimation covers a broader range of industries but a smaller number of brands, whereas my analysis focuses on a narrower set of industries while capturing a more comprehensive set of products within each industry.

B Market Coverage as a State Variable

This section provides additional evidence on the dynamics and characteristics of market coverage. Section B.1 documents how market coverage evolves over a product's life cycle, while Section B.2 presents descriptive statistics summarizing its level and adjustment patterns across products and industries.

B.1 Dynamics

Market coverage is defined as the number of retail chains in a DMA region where a product is available in each quarter. This measure strongly reflects consumers' valuation and perception of the product, as it captures how widely the product is sold across the United States. Products with broader coverage are typically more popular and exhibit higher demand. The estimated demand parameters on this availability variable, reported in Table 1, show that the coefficients ϕ are statistically significant and positive for both the cereal and yogurt industries. Specifically, a 1 percent increase in market coverage is associated with a 0.9 percentage-point increase in market share. The results highlight the importance of explicitly incorporating market coverage into the demand structure.

Furthermore, market coverage decisions are important not only for demand but also for firms' dynamic behavior. This variable is crucial for capturing producers' strategic choices and explains a substantial portion of their life-cycle dynamics. Table B1 reports the median number of stores (chain$\times$region level) following product entry, by cohort. The results show that entrants typically begin with smaller market coverage and expand it as they survive and mature. Similarly, Table B2 shows that before exit, products gradually reduce their market coverage over time. Together, these patterns confirm that market coverage decisions are inherently dynamic, influencing both demand and firms' broader strategic adjustments over the product life cycle.

Table B1: Market Coverage after Entry

	Market coverage
Entry	2
Entry + 1qt	3
Entry + 2qt	4
Entry + 3qt	5
Entry + 2yr	6
Entry + 4yr	7
Entry + 6yr	7
Entry + 8yr	8
Entry + 10yr	8

Notes: This table illustrates the median number of stores (chain×region level) in which a product is sold, by cohort shown in the first column. The statistics are based on the pooled sample across all industries.

Table B2: Market Coverage before Exit

	Market coverage
Exit – 10yr	22
Exit – 8yr	20
Exit – 6yr	12
Exit – 4yr	11
Exit – 2yr	6
Exit – 3qt	3
Exit – 2qt	2
Exit – 1qt	2
Exit	1

Notes: This table reports the median number of stores (chain×region level) in which a product is sold, by cohort shown in the first column. The statistics are based on the pooled sample across all industries.

B.2 Descriptive Statistics

Table B3 presents descriptive statistics for market coverage and its adjustments. The first row reports the summary statistics of market coverage, pooled across all industries and quarters. The average number of stores (defined at the chain×region level) is relatively large, at 121. However, the 25th percentile is only 2, and the median is about 10. Even the 75th percentile is just 92 stores, indicating that market coverage is highly right-skewed. Therefore, the empirical analysis takes the logarithm of market coverage to account for this skewness.

The second and third rows summarize adjustment decisions related to changes in market coverage. On average, when producers expand coverage, they add 19 stores, while the median increase is only 3. Conversely, when producers reduce coverage, the median decline is 4 stores, and the 25th percentile (representing larger contractions) shows a decrease of 12 stores. Overall, these summary statistics indicate that adjustment decisions are infrequent in the data, and when they occur, the magnitude of change is generally small.

Considering that approximately 34% of observations show no adjustment (last row), these patterns suggest that changes in market coverage are costly. Accordingly, the dynamic model explicitly incorporates adjustment costs to capture this behavior. Consistent with the data, the estimated dynamic parameters in Table 2 indicate high fixed and adjustment costs in the yogurt and cereal industries, reflecting the difficulty and expense associated with modifying market coverage decisions. For the adjustment parameters of other industries, see Section F.

Table B3: Summary Statistics for Market Coverage Decisions

	Mean	P25	P50	P75	SD
Market coverage	121.93	2	10	92	245.60
Increase in market coverage	19.00	1	3	9	58.29
Decrease in market coverage	–16.51	–12	–4	–1	41.41
No adjustment			34%		

Notes: The table reports summary statistics for market coverage decisions. Market coverage is defined as the number of stores (chain×region level) in which a product is sold. The second and third rows describe the magnitude of increases and decreases in market coverage, conditional on adjustment. The last row shows the share of observations with no adjustment. All statistics are based on the pooled sample across all industries.

C Details on the Estimation

C.1 Instrumental Variable Construction

The construction of instrumental variables follows the approaches of Allcott, Lockwood, and Taubinsky (2019) and Conlon and Gortmaker (2023), incorporating the local unemployment rate as in Backus, Conlon, and Sinkinson (2021). The key idea is to exploit variation in product prices across local markets within the same retail chain, which arises from uniform national pricing policies (DellaVigna and Gentzkow, 2019). Under such pricing, local demand shocks are uncorrelated with price changes in other regions, allowing these cross-market prices to serve as valid instruments for the endogenous local price (Hausman, 1996).

I first compute the unweighted national average price for each product and period. Next, I construct a Hausman-type instrument by measuring deviations of local prices from this national mean. For each brand j within chain c and region r, the regional price deviation is defined as

$$\tilde{p}_{jcrt} = p_{jcrt} - \bar{p}_{jt}, \tag{C.1}$$

where $\bar{p}_{jt}$ denotes the unweighted national average price of brand j at time t. To isolate exogenous variation, I exclude the own-region observation and compute the unweighted average of these deviations across all other regions within the same chain:

$$p^{IV}_{jcrt} = \frac{1}{R_{ct} - 1} \sum_{r' \neq r} \tilde{p}_{jcr't}, \tag{C.2}$$

where R_{ct} is the number of regions in chain c observed at time t. If a product is not sold in any other region within the same chain, the instrument is set to zero.

Using this instrument, I estimate the following first-stage regression:

$$p_{j\ell t} = \alpha_0 + \alpha_1 p^{IV}_{j\ell t} + \alpha_2 n_{j\ell t} + \mu_j + \lambda_{\ell t} + \varepsilon_{j\ell t}, \tag{C.3}$$

where $n_{j\ell t}$ denotes logged market coverage (included as a control), μ_j is a product fixed effect, and $\lambda_{\ell t}$ is a market fixed effect at the chain×region×quarter level. The fitted values $\hat{p}_{j\ell t}$ from this regression are then used to construct the differentiated instrument of Gandhi and Houde (2019).

This differentiated instrumental variable measures the relative price distance across products within each local market. Specifically, for product j in market ℓ, I compute

$$p_{j\ell t}^{\text{IV,diff}} = \sum_{k \neq j \in J_{\ell t}} \left(\hat{p}_{j\ell t} - \hat{p}_{k\ell t}\right)^2 , \tag{C.4}$$

where $\hat{p}_{j\ell t}$ denotes the predicted price from the first-stage regression. This quadratic instrument performs substantially better in capturing exogenous price variation within markets.

I also construct an analogous quadratic version of the differentiated instrument based on logged market coverage n_{jt}, defined as

$$n_{j\ell t}^{\text{IV,diff}} = \sum_{k \neq j \in J_{\ell t}} \left(n_{j\ell t} - n_{k\ell t}\right)^2 . \tag{C.5}$$

Although N_{jt} is fixed for each product at time t, within a given market ℓ the coverage levels N_{kt} for other products ($k \neq j$) differ, reflecting the differentiated nature of product availability.

Together, these three IV variables provide substantial first-stage F-statistics for each industry, and the resulting demand parameters are all precisely estimated.

C.2 Price Elasticity of Demand

Based on the estimated demand parameters in Section 4 and the recovered product appeal, I calculate the price elasticity of demand.

Using the individual demand function (5) and the aggregate demand function (8), the analytical form of the demand derivative with respect to the own price is

$$\frac{\partial D_{j\ell t}}{\partial p_{j\ell t}} \frac{p_{j\ell t}}{D_{j\ell t}} = \int \alpha_i d_{ij\ell t}(d_{ij\ell t} - 1)\, dF(\alpha_i) \frac{p_{j\ell t}}{\int d_{ij\ell t}\, dF(\alpha_i)} \tag{C.6}$$

which is typically evaluated numerically using Gauss–Hermite quadrature.

Approximating the integral with discrete nodes $\{\alpha_i, w_i\}_{i=1}^{I}$, it can be expressed as

$$\int d_{ij\ell t} dF(\alpha_i) \approx \frac{1}{\sqrt{\pi}} \sum_{i=1}^{I} w_i \alpha_i d_{ij\ell t} \tag{C.7}$$

and

$$\int \alpha_i d_{ij\ell t}(d_{ij\ell t} - 1)\, dF(\alpha_i) \approx \frac{1}{\sqrt{\pi}} \sum_{i=1}^{I} w_i \alpha_i d_{ij\ell t}(d_{ij\ell t} - 1). \tag{C.8}$$

I set the order of the quadrature to $I = 11$. This calculation is performed at the local-market (ℓ) level. The computation is feasible because the parameters of the distribution $(\alpha, \sigma_\alpha^2)$ as well as the availability parameter ϕ are identified. Prices and market coverage are taken directly from the data, and the recovered product appeal from the estimation is used in the calculation. For the main analysis, this recovered elasticity at the product-level is used for comparative analysis. For the cross-industry analysis, I use the median price elasticity of demand within each industry.

C.3 National-Level Aggregation

The aggregation of product appeal to the national level is defined as

$$\xi_{jt} = \sum_{\ell \in \overline{L}} \omega_{j\ell t} \left(\xi_{j\ell t} + \xi_j \right), \tag{C.9}$$

where ℓ indexes local markets (defined by chain and region), and $\omega_{j\ell t}$ denotes the corresponding quantity weight. ξ_j represents the estimated product fixed effect from the demand estimation, capturing time-invariant characteristics that influence demand beyond time-varying components.

The demand estimation includes two sets of fixed effects: a brand fixed effect and a market fixed effect at the chain$\times$region$\times$quarter level. In some cases, products with limited observations do not provide sufficient variation to identify these fixed effects directly. For these cases, the corresponding group-level effects are interpolated to retain niche products that would otherwise be excluded from the sample.

Prices are aggregated in a similar manner:

$$p_{jt} = \sum_{\ell \in \overline{L}} \omega_{j\ell t} \, p_{j\ell t}. \tag{C.10}$$

The total number of local markets, $\overline{L}$, is fixed over time. Since a product may not be sold in all local markets, the corresponding weight $\omega_{j\ell t}$ is set to zero when sales are not observed in market ℓ at time t.

C.4 New Entrant State

The entrant-specific product appeal (ξ) distribution is obtained from the data for each industry. Figure C2 presents the pooled distribution of entrant appeal across industries, which I assume to be time-invariant, as there is no evidence of systematic time variation in the data. After standardizing appeal within each industry, the distribution exhibits a well-behaved bell shape without a distinct mode. Therefore, I do not assume that entrants start from a specific level of product appeal; instead, each entrant draws its initial ξ from this empirically derived distribution.

This contrasts with the market coverage state of entrants. As shown in Section B.1, entrants typically begin with relatively small market coverage and expand as they survive and remain active in the economy. Given that market coverage equals one in many observed cases, I set the entrant's initial market coverage at this level in the model.

For the state-space construction in the dynamic model, I discretize the support of ξ using the empirical distribution. The minimum and maximum correspond to the 1st and 99th percentiles within each industry, and the grid points are set to be equally spaced within this range.

Figure C2: Product Appeal Distribution of Entrants

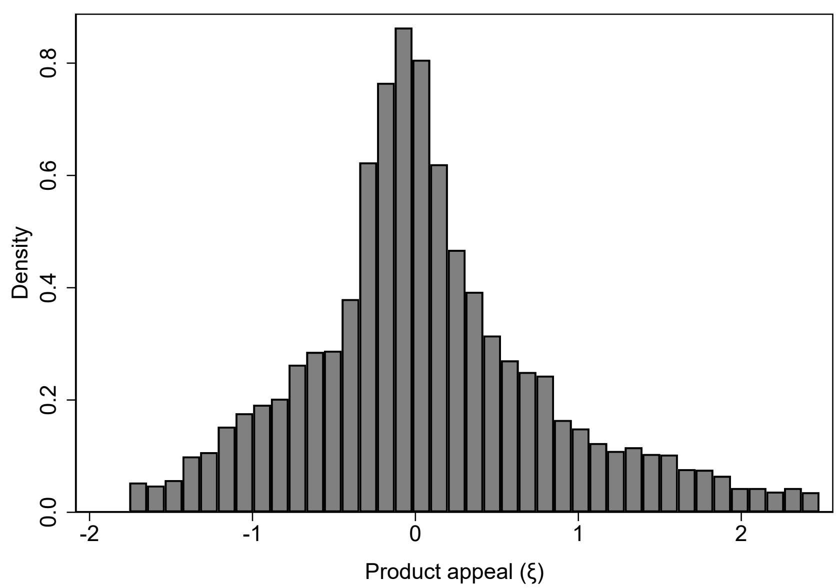

Note: The figure plots the distribution of recovered product appeal (ξ) for newly introduced products, pooled across all industries in the sample. To make the distribution comparable across industries, the appeal values are standardized within each industry before pooling.

D Price Function Specification

To incorporate prices into the dynamic estimation (i.e., quarterly prices for each product), I require data for every state variable. However, some state variables are missing in the data. As explained in Section B, market coverage is highly right-skewed, resulting in sparse observations for states in the upper tail of the distribution. To address this issue, I use a second-order nonparametric regression to impute prices for missing states and construct a complete price function for the dynamic model.

Accounting for the relationship between prices and the recovered state variables in the data helps illustrate why this approach is relevant. First, product appeal is positively associated with price, as shown in Figure D3. This relationship aligns with the idea that products with higher appeal command higher prices. It can also be interpreted as a reflection of higher quality, even though product appeal in this paper captures a broader set of demand-side factors beyond quality. Similarly, Figure D4 shows that products with greater market coverage tend to have lower prices, consistent with economies of scale. These patterns indicate that the two state variables, product appeal and market coverage, jointly explain a significant share of price variation. Accordingly, I model the price function nonparametrically as a flexible function of these state variables over time.

Specifically, I estimate

$$p_{jt} = g\left(\xi_{jt}, \log N_{jt}, t\right) + \varepsilon_{jt}, \tag{D.1}$$

where $g(\cdot)$ is a nonparametric function of the state variables and time. For implementation, I approximate $g(\cdot)$ with a second-order polynomial. Let $n_{jt} \equiv \log N_{jt}$. The estimating equation is

$$\begin{aligned} p_{jt} = \beta_0 &+ \beta_1 \xi_{jt} + \beta_2 n_{jt} + \beta_3 t \\ &+ \beta_4 \xi_{jt}^2 + \beta_5 n_{jt}^2 + \beta_6 t^2 \\ &+ \beta_7 \xi_{jt} n_{jt} + \beta_8 \xi_{jt} t + \beta_9 n_{jt} t + \varepsilon_{jt}. \end{aligned} \tag{D.2}$$

This specification allows prices to vary over time while flexibly capturing how they co-move with the observed state transitions in the data. In the model, the optimal price can then be written as

$$p^*_{jt} = \hat{g}\left(\xi_{jt}, \log N_{jt}, t\right). \tag{D.3}$$

Figure D3: Average Price by Product Appeal Decile

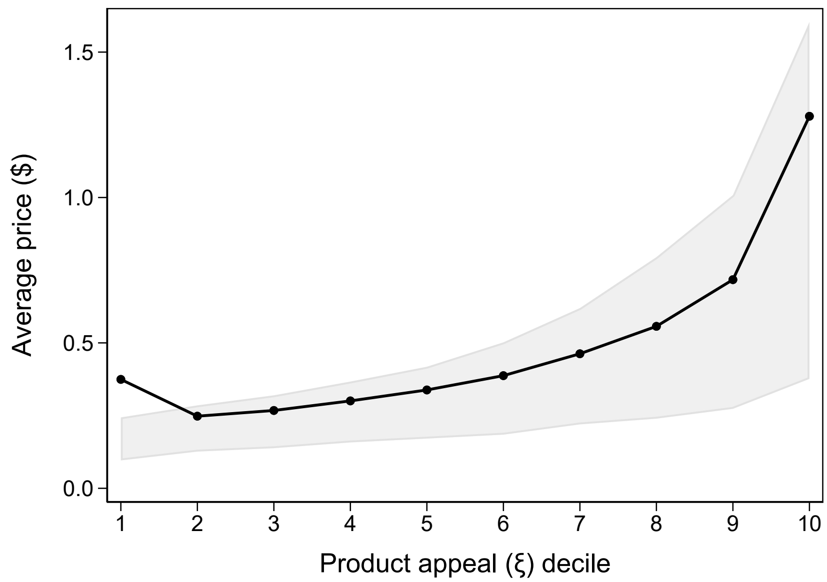

Note: The figure plots the average product price by deciles of product appeal (ξ), pooled across all industries in the sample. The solid line represents the mean price within each decile, and the shaded area indicates the interquartile range (25th-75th percentiles).

Figure D4: Average Price by Market Coverage Decile

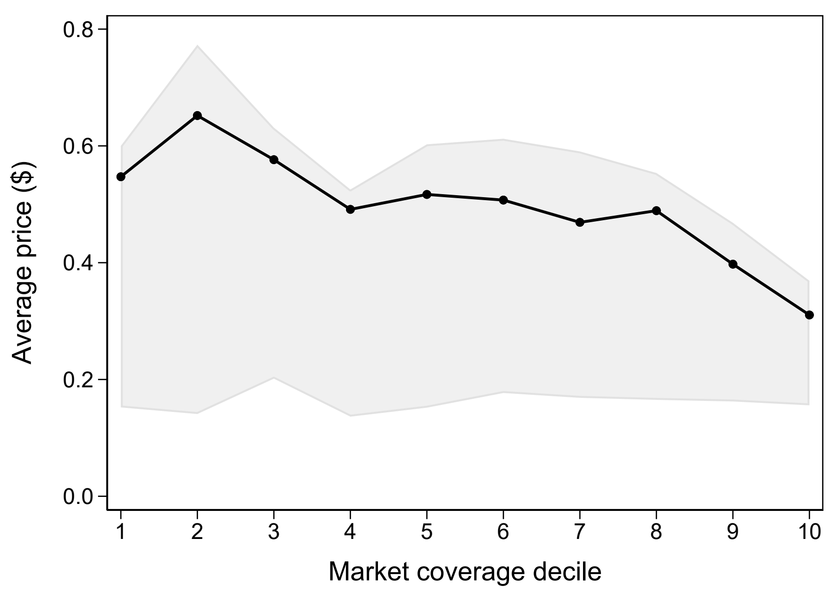

Note: The figure plots the average product price by deciles of market coverage ($\log N_{jt}$), pooled across all industries in the sample. The solid line represents the mean price within each decile, and the shaded area indicates the interquartile range (25th-75th percentiles).

E More on Welfare Analysis

The welfare analysis begins by comparing economies with and without aggregate shocks. For each industry, I simulate the model 500 times to account for uncertainty in the policy functions, which are expressed as probabilities. Random draws across simulations provide reasonable confidence intervals for welfare outcomes. In each phase's initial period, the number of products is initialized using the observed data. For comparability, each simulation uses a distinct random seed that is shared across the shock and no-shock cases, so that the only difference arises from the realization of aggregate shocks.

Because the model incorporates non-stationarity, I begin by evaluating welfare separately for each period, discounting all values to the present using a social discount factor of $\beta_s = 0.999$. Figure E5 shows the time-varying gap in social welfare between the two economies. The dynamics of welfare decline closely follow the path of the calibrated aggregate demand shock.

Most of the welfare loss originates from the decline in consumer surplus. Figure E6 decomposes the welfare gap into consumer surplus (Panel A) and producer surplus (Panel B). The decline in consumer surplus is substantial and statistically significant, whereas producer surplus increases slightly at the onset of the shock. This pattern arises because the shock induces a strong rise in product exit and a decline in product entry. Although firm profits eventually fall, the immediate effect of lower entry costs and positive exit values generates a small short-run gain in producer surplus.

Figure E5: Social Welfare Difference over Time

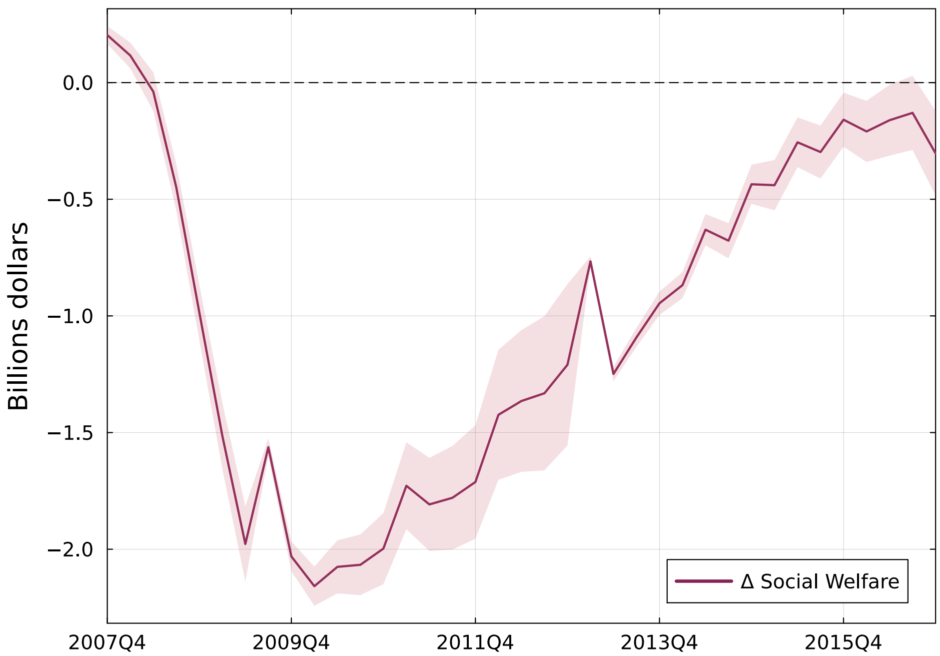

Note: The figure plots the time-varying difference in total social welfare between the economies with and without aggregate shocks. Shaded areas indicate 95% confidence intervals based on 500 Monte Carlo simulations. Welfare is expressed in billions of dollars, discounted using $\beta_s = 0.999$.

Figure E6: Decomposition: Welfare Gap over Time

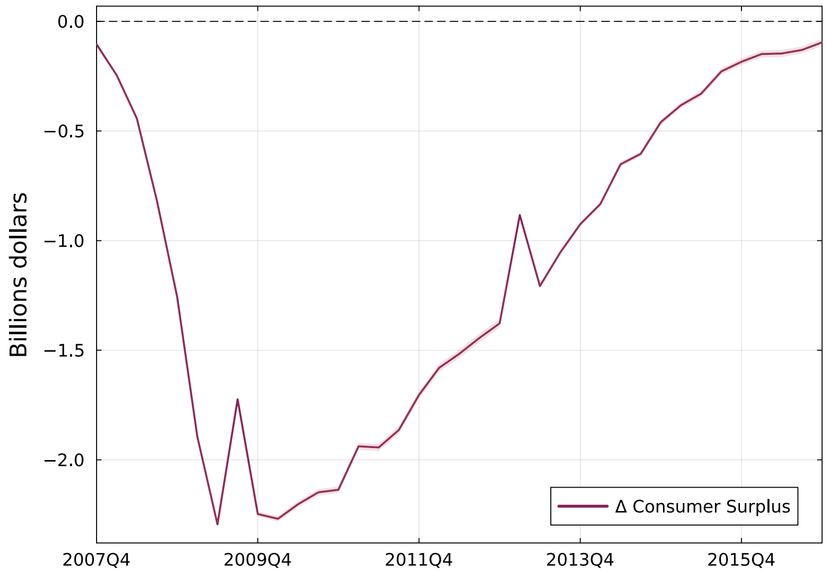

(a) Welfare Gap in Consumer Surplus

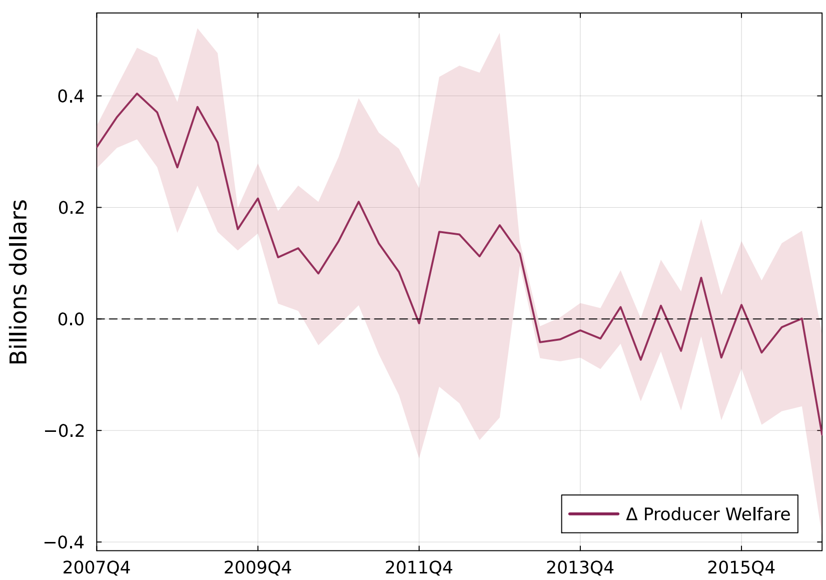

(b) Welfare Gap in Producer Surplus

Note: The figure decomposes the social welfare gap into consumer surplus (Panel a) and producer surplus (Panel b) between economies with and without aggregate shocks. Shaded areas indicate 95% confidence intervals based on 500 Monte Carlo simulations. Welfare is expressed in billions of dollars, discounted by the social discount factor $\beta_s = 0.999$.

F Additional Tables and Figures

Section F.1 reports the estimated demand parameters for the remaining industries. Section F.2 presents the estimated dynamic parameters for the remaining industries and provides full, phase-specific estimates, including later-period results for cereals and yogurt. Finally, Section F.3 documents model fit across all industries and shows that the estimated model tracks the key patterns in the data.

F.1 Demand Estimation

Table F4: Estimates of Demand Parameters: Other Industries

Parameters	Notation	Canned soup	Pasta (spaghetti)	Cookies	Jams	Honey	Peanut butter	Bottled water
Price sensitivity (mean)	α	-15.69	-24.01	-26.28	-33.75	-16.07	-29.00	-88.49
		(4.90)	(8.70)	(4.49)	(11.94)	(2.38)	(3.91)	(3.37)
Price sensitivity (variance)	σ_α^2	11.44	47.20	12.21	27.26	0.64	37.33	44.40
		(5.57)	(23.67)	(2.62)	(11.08)	-0.10	(5.91)	(1.62)
Product availability	ϕ	1.00	0.79	0.75	0.20	0.51	0.61	0.82
		(0.09)	(0.04)	(0.04)	(0.08)	(0.03)	(0.01)	(0.01)
Autocorrelation of product appeal		0.96	0.92	0.53	0.97	0.96	0.93	0.96
Median price elasticity of demand		-1.49	-2.05	-6.29	-3.31	-4.63	-3.83	-3.68

Table F4: Estimates of Demand Parameters: Other Industries (continued)

Parameters	Notation	Chocolate	Grated cheese	Specialty soap	Hand cleaners	Tooth cleaners	Deodorant
Price sensitivity (mean)	α	-20.73	-8.15	-12.24	-0.93	-10.55	-11.86
		(0.79)	(1.28)	(1.98)	(0.13)	(2.89)	(0.47)
Price sensitivity (variance)	σ_α^2	3.61	3.04	1.21	0.02	7.28	2.99
		(0.19)	(1.01)	-0.30	0.00	(3.35)	(0.13)
Product availability	ϕ	0.84	0.39	1.08	0.80	1.12	1.55
		(0.01)	(0.01)	(0.04)	(0.01)	(0.10)	(0.02)
Autocorrelation of product appeal		0.76	0.94	0.92	0.87	0.93	0.93
Median price elasticity of demand		-6.93	-2.95	-2.90	-0.28	-3.98	-10.04

Notes: Robust standard errors are in parentheses. Hand cleaners imply Hand cleaners and sanitizers.

F.2 Estimating Dynamic Model

Table F5: Estimates of Dynamic Parameters: Cereal and Yogurt-Full Phase

Parameters	Notation	Cereal		Yogurt	
		2009Q3-2012Q4	2013Q1-2016Q4	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	747.3	500.63	394.73	499.39
		(270.41)	(112.71)	(128.25)	(153.91)
Mean exit value ($mn)	γ^{ex}	107.88	63.73	40.84	52.31
		(40.3)	(14.92)	(14.02)	(16.65)
Fixed cost of market coverage adjustment ($mn)	γ_0	2.66	2.81	2.32	2.69
		(0.04)	(0.04)	(0.09)	(0.08)
Linear cost of increasing market coverage ($k)	γ_1^+	52.32	58.74	81.9	85.22
		(2.01)	(1.81)	(4.6)	(3.39)
Linear cost of decreasing market coverage ($k)	γ_1^-	-18.34	-9.08	-36.53	-6.36
		(3.46)	(3.22)	(9)	(6.4)
Convex cost of market coverage adjustment	γ_2	1.47	0.88	1.27	0.83
		(23.98)	(20.84	(69.93	(39.37
Logit shock scale	ψ	0.69	0.83	0.36	0.95
		(0.18)	(0.08)	(0.12)	(0.09)
Sunk entry cost ($mn)		639.42	436.90	353.89	447.08

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase

Parameters	Notation	Canned soup		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	800.81	533.81	276.11
		(241.96)	(107.35)	(42.35)
Mean exit value ($mn)	γ^{ex}	32.22	28.37	24.29
		(6.18)	(3.84)	(2.64)
Fixed cost of market coverage adjustment ($mn)	γ_0	2.07	2.35	2.29
		(0.13)	(0.09)	(0.08)
Linear cost of increasing market coverage ($k)	γ_1^+	200.46	197	228.86
		(13.41)	(6.52)	(5.63)
Linear cost of decreasing market coverage ($k)	γ_1^-	143.38	44.7	53.48
		(27.7)	(13.57)	(12.27)
Convex cost of market coverage adjustment	γ_2	2397.32	512.25	491.62
		(428.34)	(98.27	(83.97
Logit shock scale	ψ	0.54	0.01	0.16
		(0.16)	(0.03)	(0.05)
Sunk entry cost ($mn)		768.59	505.44	251.82

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase (continued)

Parameters	Notation	Pasta (spaghetti)		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	911.64	350.59	285.21
		(669.89)	(120.63)	(67.09)
Mean exit value ($mn)	γ^{ex}	63.74	28.95	22.36
		(48.2)	(10.37)	(5.34)
Fixed cost of market coverage adjustment ($mn)	γ_0	1.89	2.43	2.4
		(0.16)	(0.13)	(0.11)
Linear cost of increasing market coverage ($k)	γ_1^+	92.52	68.69	63.61
		(12.77)	(8.51)	(7.29)
Linear cost of decreasing market coverage ($k)	γ_1^-	-112.91	-110.76	-112.84
		(24.34)	(17.93)	(15.62)
Convex cost of market coverage adjustment	γ_2	5.16	3.63	3.18
		(394.82)	(248.1	(202.56
Logit shock scale	ψ	1.27	0.94	1.02
		(0.07)	(0.05)	(0.02)
Sunk entry cost ($mn)		847.90	321.64	262.85

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase (continued)

Parameters	Notation	Cookies		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	881.91	762.42	695.54
		(290.57)	(309.47)	(273.06)
Mean exit value ($mn)	γ^{ex}	141.27	161.86	141.65
		(47.46)	(66.97)	(56.86)
Fixed cost of market coverage adjustment ($mn)	γ_0	2.67	2.77	2.86
		(0.05)	(0.04)	(0.04)
Linear cost of increasing market coverage ($k)	γ_1^+	59.52	50.67	56.36
		(2.22)	(1.6)	(1.48)
Linear cost of decreasing market coverage ($k)	γ_1^-	-14.84	-18.12	-12.74
		(4.47)	(3.29)	(3.08)
Convex cost of market coverage adjustment	γ_2	0.57	0.53	0.59
		(29.85)	(21.92	(20.35
Logit shock scale	ψ	0.16	1.08	0.4
		(0.65)	(0.14)	(0.39)
Sunk entry cost ($mn)		740.64	600.56	553.89

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase (continued)

Parameters	Notation	Jams		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	414	581.77	175.35
		(176.52)	(305.25)	(40.35)
Mean exit value ($mn)	γ^{ex}	16.62	33.53	18.08
		(7.1)	(18.48)	(4.44)
Fixed cost of market coverage adjustment ($mn)	γ_0	1.11	1.97	2.25
		(0.38)	(0.3)	(0.25)
Linear cost of increasing market coverage ($k)	γ_1^+	170.58	110.89	114.45
		(35.76)	(24.42)	(18.47)
Linear cost of decreasing market coverage ($k)	γ_1^-	-254.03	-189.95	-142.37
		(78.48)	(55.14)	(43.76)
Convex cost of market coverage adjustment	γ_2	0.23	38.38	0.03
		(2085.68)	(1135.38	(800.1
Logit shock scale	ψ	1	0.82	0.79
		(0.08)	(0.1)	(0.05)
Sunk entry cost ($mn)		397.38	548.24	157.27

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase (continued)

Parameters	Notation	Honey		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	323.61	263.92	137.92
		(85.31)	(55.86)	(19.02)
Mean exit value ($mn)	γ^{ex}	14.89	20.85	15.4
		(3.6)	(4.52)	(2.16)
Fixed cost of market coverage adjustment ($mn)	γ_0	2.37	2.36	2.27
		(0.23)	(0.17)	(0.16)
Linear cost of increasing market coverage ($k)	γ_1^+	91.94	95.5	85.77
		(15.11)	(10.96)	(10.64)
Linear cost of decreasing market coverage ($k)	γ_1^-	-103.16	-105.87	-140.44
		(35.53)	(26.21)	(25.33)
Convex cost of market coverage adjustment	γ_2	0.01	0	5.34
		(511.03)	(372	(374.29
Logit shock scale	ψ	1.01	0.96	0.8
		(0.03)	(0.05)	(0.08)
Sunk entry cost ($mn)		308.72	243.07	122.52

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase (continued)

Parameters	Notation	Peanut butter		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	885.99	934.74	335.25
		(490.87)	(586.05)	(94.61)
Mean exit value ($mn)	γ^{ex}	19.9	58.75	20.56
		(10.5)	(38.49)	(6.02)
Fixed cost of market coverage adjustment ($mn)	γ_0	2.8	2.39	2.45
		(0.22)	(0.13)	(0.11)
Linear cost of increasing market coverage ($k)	γ_1^+	32.23	39.13	39.3
		(9.97)	(7.58)	(5.76)
Linear cost of decreasing market coverage ($k)	γ_1^-	-71.84	-88.1	-91.87
		(21.11)	(14.4)	(12.39)
Convex cost of market coverage adjustment	γ_2	3.01	1.47	1.33
		(171.12)	(141.2	(115.19
Logit shock scale	ψ	1.19	0.16	1.05
		(0.04)	(0.43)	(0.03)
Sunk entry cost ($mn)		866.09	875.99	314.69

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase (continued)

Parameters	Notation	Bottled water		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	862.59	898.02	875.58
		(480.17)	(336.64)	(275.89)
Mean exit value ($mn)	γ^{ex}	121.58	146.43	143.18
		(69.28)	(56.07)	(46.17)
Fixed cost of market coverage adjustment ($mn)	γ_0	2.59	2.46	2.38
		(0.1)	(0.07)	(0.06)
Linear cost of increasing market coverage ($k)	γ_1^+	85.69	82.43	89.51
		(5.23)	(3.99)	(3.63)
Linear cost of decreasing market coverage ($k)	γ_1^-	-37.32	-54.42	-54.73
		(10.96)	(8.42)	(7.5)
Convex cost of market coverage adjustment	γ_2	1.65	1.18	1.34
		(111.69)	(93.72	(85.5
Logit shock scale	ψ	1.05	0.82	1.53
		(0.18)	(0.4)	(0.06)
Sunk entry cost ($mn)		741.01	751.59	732.40

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase (continued)

Parameters	Notation	Chocolate		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	492.92	866.82	894.81
		(84.19)	(238.7)	(182.32)
Mean exit value ($mn)	γ^{ex}	61.36	178.68	155.57
		(11.03)	(50.19)	(32.39)
Fixed cost of market coverage adjustment ($mn)	γ_0	2.57	2.63	2.83
		(0.05)	(0.04)	(0.04)
Linear cost of increasing market coverage ($k)	γ_1^+	30.76	25.89	22.83
		(2.23)	(1.64)	(1.36)
Linear cost of decreasing market coverage ($k)	γ_1^-	-39.75	-45.75	-36.97
		(4.15)	(3)	(2.64)
Convex cost of market coverage adjustment	γ_2	0.77	0.63	0.51
		(28.58)	(21.37	(16.62
Logit shock scale	ψ	0.45	1.76	1.29
		(0.4)	(0.16)	(0.12)
Sunk entry cost ($mn)		431.56	688.14	739.24

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase (continued)

Parameters	Notation	Grated cheese		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	527.59	251.39	194.13
		(427.58)	(73.42)	(45.12)
Mean exit value ($mn)	γ^{ex}	34.86	17.24	14.56
		(28.99)	(4.89)	(3.2)
Fixed cost of market coverage adjustment ($mn)	γ_0	2.7	2.1	1.81
		(0.29)	(0.2)	(0.19)
Linear cost of increasing market coverage ($k)	γ_1^+	60.83	73.07	80.29
		(15.86)	(12.54)	(13.55)
Linear cost of decreasing market coverage ($k)	γ_1^-	-69.86	-94.42	-154.34
		(35.14)	(28.3)	(29.43)
Convex cost of market coverage adjustment	γ_2	0	1.65	2.31
		(370.7)	(360.47	(439.6
Logit shock scale	ψ	1.02	1.09	0.97
		(0.11)	(0.02)	(0.03)
Sunk entry cost ($mn)		492.73	234.15	179.57

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase (continued)

Parameters	Notation	Specialty soap		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	755.63	525.36	500.13
		(428.5)	(199.12)	(185.73)
Mean exit value ($mn)	γ^{ex}	109.55	74.17	76.24
		(63.94)	(29.17)	(29.28)
Fixed cost of market coverage adjustment ($mn)	γ_0	1.96	2.38	2.49
		(0.06)	(0.04)	(0.04)
Linear cost of increasing market coverage ($k)	γ_1^+	83.92	70.19	58.09
		(3.07)	(2.02)	(1.59)
Linear cost of decreasing market coverage ($k)	γ_1^-	-23.77	-22.27	-22.35
		(5.53)	(4.27)	(3.51)
Convex cost of market coverage adjustment	γ_2	1.66	1.2	0.82
		(36)	(25.14	(20.24
Logit shock scale	ψ	0.73	0.34	0.89
		(0.4)	(0.14)	(0.08)
Sunk entry cost ($mn)		646.08	451.19	423.89

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase (continued)

Parameters	Notation	Hand cleaners and sanitizers		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	797.3	788.39	985.84
		(908.14)	(366.92)	(200.59)
Mean exit value ($mn)	γ^{ex}	89.29	55.74	53.6
		(105.38)	(26.98)	(11.13)
Fixed cost of market coverage adjustment ($mn)	γ_0	2.14	1.96	1.93
		(0.14)	(0.09)	(0.11)
Linear cost of increasing market coverage ($k)	γ_1^+	69.31	95.43	109.55
		(7.65)	(5.28)	(6.19)
Linear cost of decreasing market coverage ($k)	γ_1^-	-27.63	-29.55	-20.14
		(13.46)	(10.23)	(12.87)
Convex cost of market coverage adjustment	γ_2	0.79	0.93	1.32
		(117.9)	(98.78	(122.75
Logit shock scale	ψ	0.01	1.35	0.01
		(1.8)	(0.13)	(0.17)
Sunk entry cost ($mn)		708.01	732.65	932.24

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase (continued)

Parameters	Notation	Tooth cleaners		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	378	332.74	382.1
		(121.78)	(83.67)	(90.91)
Mean exit value ($mn)	γ^{ex}	17.31	21.51	25.72
		(5.09)	(5.43)	(6.12)
Fixed cost of market coverage adjustment ($mn)	γ_0	2.76	3.04	3.04
		(0.11)	(0.08)	(0.07)
Linear cost of increasing market coverage ($k)	γ_1^+	38.19	36.54	39.79
		(5.16)	(4.12)	(3.3)
Linear cost of decreasing market coverage ($k)	γ_1^-	-50.27	-55.11	-41.42
		(10.39)	(8.74)	(7.05)
Convex cost of market coverage adjustment	γ_2	1.46	3.72	6.64
		(88.69)	(77.37	(56.02
Logit shock scale	ψ	1.23	1.02	1.07
		(0.04)	(0.04)	(0.02)
Sunk entry cost ($mn)		360.69	311.23	356.38

Table F6: Estimates of Dynamic Parameters: Other Industries-Full Phase (continued)

Parameters	Notation	Deodorant		
		2007Q4-2009Q2	2009Q3-2012Q4	2013Q1-2016Q4
Mean entry cost ($mn)	γ^{en}	534.68	588.99	478.57
		(256.65)	(244.61)	(192.75)
Mean exit value ($mn)	γ^{ex}	26.83	35.27	40.18
		(13.36)	(15.38)	(17.09)
Fixed cost of market coverage adjustment ($mn)	γ_0	3.23	3.09	3.19
		(0.13)	(0.08)	(0.08)
Linear cost of increasing market coverage ($k)	γ_1^+	20.35	27.05	23.25
		(5.63)	(3.69)	(3.34)
Linear cost of decreasing market coverage ($k)	γ_1^-	-53.42	-49.15	-50.04
		(11.51)	(7.52)	(7.07)
Convex cost of market coverage adjustment	γ_2	1.73	1.54	1.35
		(86.89)	(57.89	(51.76
Logit shock scale	ψ	1.2	1.19	0.93
		(0.08)	(0.02)	(0.05)
Sunk entry cost ($mn)		507.85	553.72	438.39

F.3 Model Fit

Figure F7: Model Fit: Other Industries

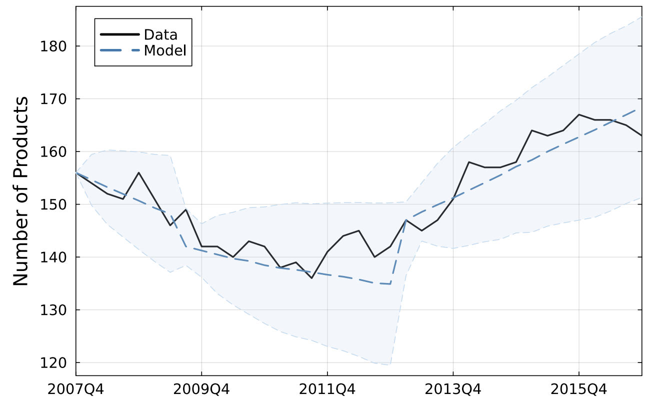

(a) Canned soup

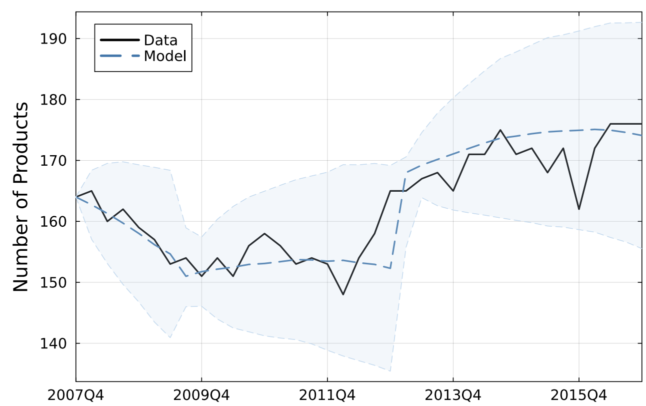

(b) Pasta (spaghetti)

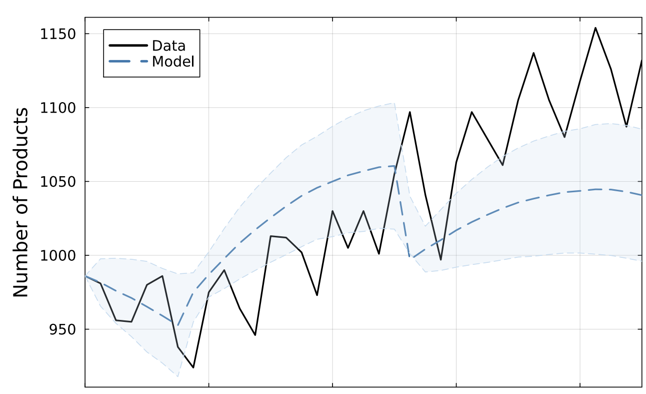

(c) Cookies

Figure F7: Model Fit: Other Industries (continued)

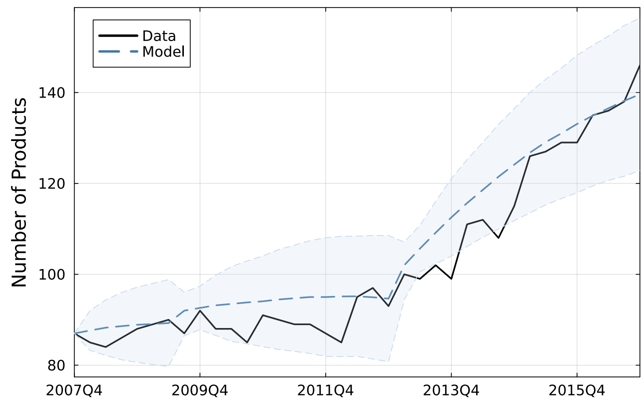

(d) Jams

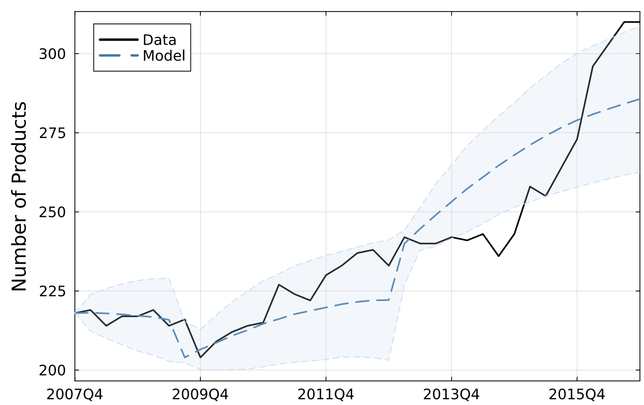

(e) Honey

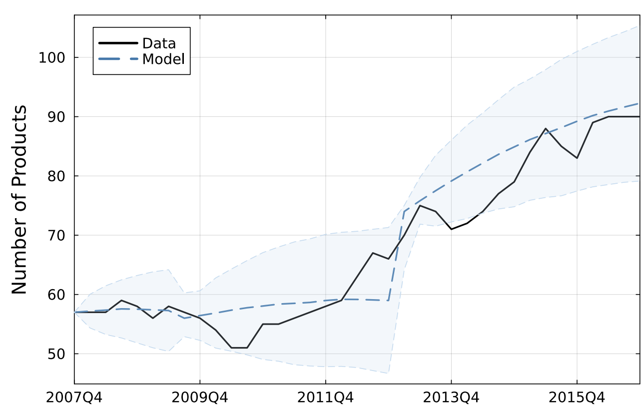

(f) Peanut Butter

Figure F7: Model Fit: Other Industries (continued)

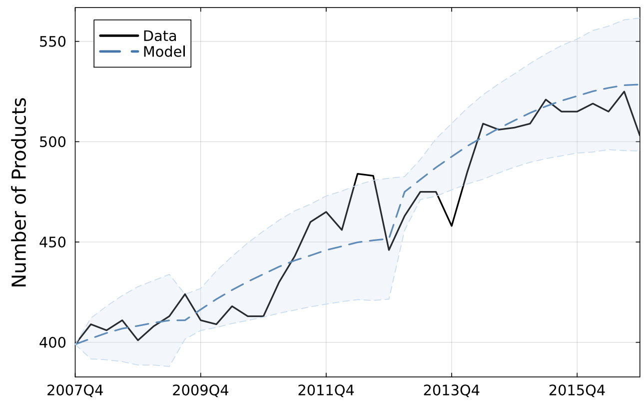

(g) Bottled water

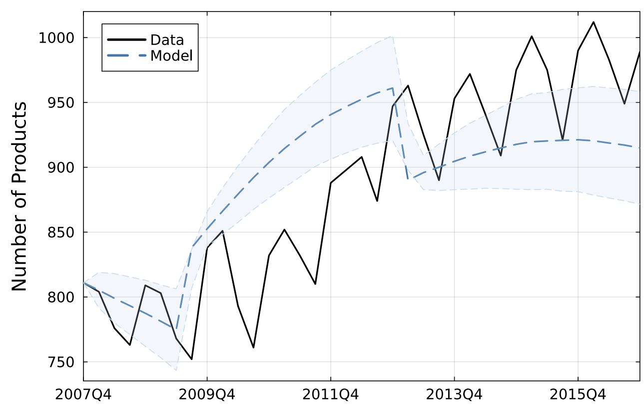

(h) Chocolate

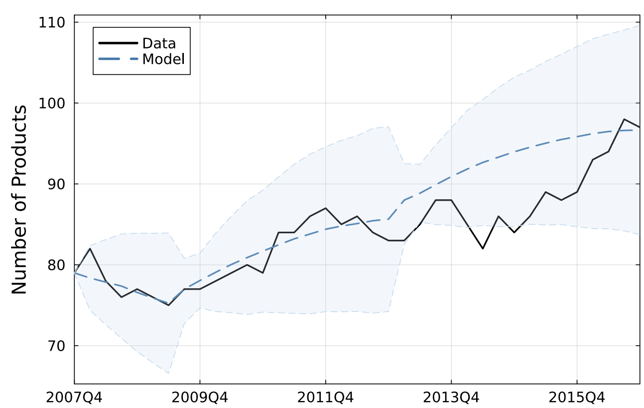

(i) Grated Cheese

Figure F7: Model Fit: Other Industries (continued)

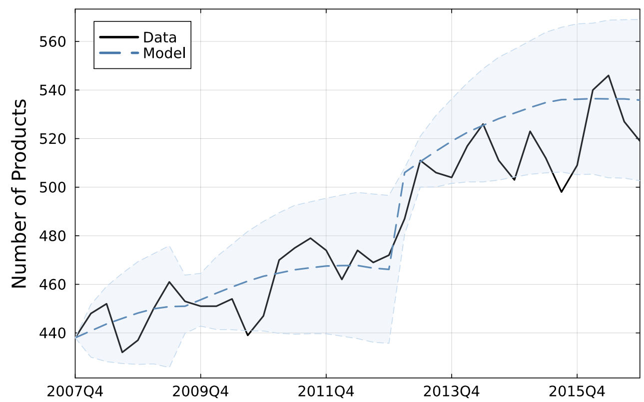

(j) Specificalty soap

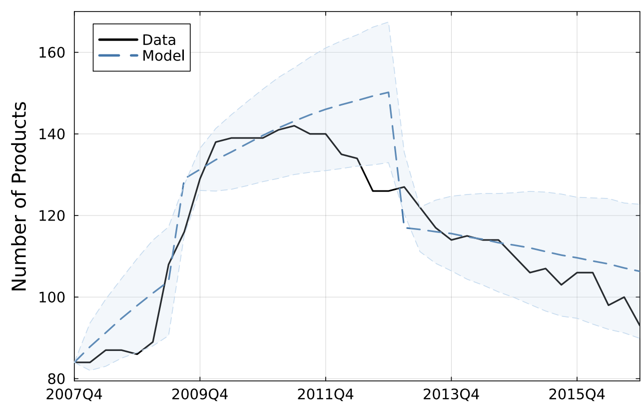

(k) Hand cleaners and sanitizers

Figure F7: Model Fit: Other Industries (continued)

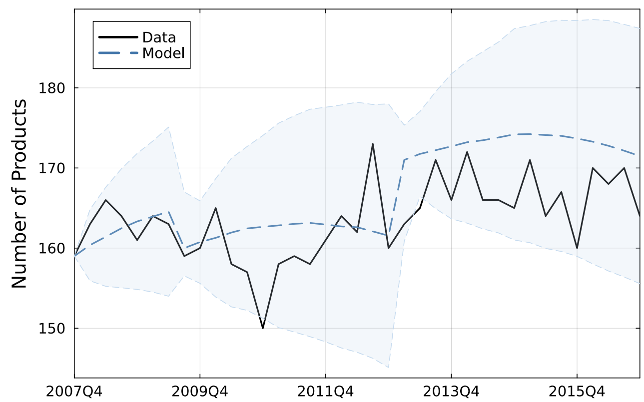

(l) Tooth cleaners

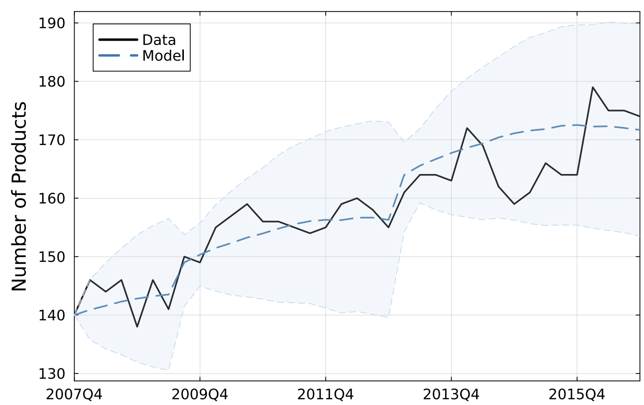

(m) Deodorant