Mergers in Markets with Consumer Inertia*

Alexander MacKay
University of Virginia†

Marc Remer
Swarthmore College‡

June 2026

Abstract

We study how consumer inertia shapes the price effects of horizontal mergers. With inertia—arising from habit formation, brand loyalty, or switching costs—firms internalize how current prices affect future demand, creating dynamic pricing incentives absent from static models. We develop an empirical oligopoly model that captures consumer inertia and dynamic pricing incentives with market-level data. Applying the model to retail gasoline markets, where we document reduced-form evidence consistent with dynamics, we estimate that 63 percent of consumers are subject to inertia. Merger simulations show that price effects vary materially with the post-merger structure, specifically, whether the merged firm consolidates brands or maintains them separately with joint pricing. Further, static merger analysis can substantially misstate price effects. In our application, a brand consolidation merger increases prices by 2.4 percent, while a static model predicts a 5.4 percent price increase because it omits the merged firm's incentive to invest in future demand.

JEL Codes: D12, D43, L13, L41, L81

Keywords: Consumer Inertia, Market Power, Dynamic Competition, Demand Estimation

---

*An earlier version of this paper was entitled "Consumer Inertia and Market Power." We are grateful to Ali Hortaçsu, Robin Lee, Nathan Miller, Aviv Nevo, Ariel Pakes, Karl Schurter, and Nathan Wilson for their comments. We thank participants at the IIOC conference and the EEA-ESEM conference, and also seminar participants at the University of Toronto, Haverford College, Swarthmore College, Harvard Business School, the University of Florida, Boston College, Bowdoin College, U.S. Department of Justice, NYU Stern, and Harvard University for their feedback and suggestions. Researchers' 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 researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

†University of Virginia, Department of Economics. Email: mackay@virginia.edu.

‡Swarthmore College, Department of Economics. Email: mremer1@swarthmore.edu.

1 Introduction

Consumers are often more likely to buy a product if they have purchased it previously. This tendency may reflect consumer inertia, which can arise from habit formation, brand loyalty, or switching costs. In response to consumer inertia, profit-maximizing firms will internalize the effect of their current price on demand in future periods.

In this paper, we study the influence of consumer inertia on competition and pricing in oligopoly settings. We focus in particular on horizontal mergers in markets characterized by consumer inertia. To perform the analysis, we implement a model of inertia where a portion of consumers develop a product-specific preference, or *affiliation*, for the product they purchased most recently. We then use the model to evaluate how consumer inertia affects equilibrium prices and merger outcomes.

We make two primary contributions related to the analysis of mergers. First, we show that the effects of mergers on equilibrium prices depend on the presence of inertia. Using our empirical model that allows for consumer inertia, we show that the predictions of a static model can diverge meaningfully from the true effects, primarily due to misspecified first-order conditions. Second, we highlight an important feature of merger implementation. After a merger, a firm may decide to maintain separate brands, which we call a *joint pricing* merger, or fold the products into a single brand, i.e., a *brand consolidation* merger. We show that this decision, which may not matter in static models, can have substantive implications in the presence of inertia. Both outcomes occur in practice. For example, in the airline industry, mergers typically result in one of the brands being eliminated.[1] Conversely, with consumer product mergers, the acquiring firm often maintains both existing and acquired brands.[2] In retail gasoline, both types of mergers have occurred, with BP eliminating the Amoco brand after its 1998 acquisition, while Exxon maintained the Mobil brand for service stations after its acquisition in 1999.[3]

With theoretical analysis and empirical merger simulations, we illustrate how consumer dynamics can influence post-merger outcomes. In a static setting, distinguishing between merger implementation types may not be relevant; we show that mergers of either type deliver the exact same outcome for symmetric firms with logit demand. However, in the presence of consumer inertia, differences in dynamic pricing incentives can lead these two merger types to deliver materially different effects relative to each other and to the static model. Accounting for these differences is relevant for the decisions of firms and competition authorities, which typically employ static empirical models to estimate consumer demand and simulate counterfactual post-merger prices. Such static models could, for instance, substantially overpredict or

[1] Historical examples include Delta-Northwest, United-Continental, and Southwest-AirTran mergers, leaving Delta, United, and Southwest, respectively.

[2] Coca-Cola launched Gold Peak Tea in 2006, acquired Honest Tea in 2011, and acquired Peace Tea in 2015. It maintained these three separate ready-to-drink iced tea brands through 2022.

[3] Interestingly, BP reintroduced the Amoco brand for retail stations in 2017.

underpredict the price effects of a merger depending on the magnitude of inertia and whether the merged firm plans to consolidate brands. Thus, failing to account for consumer inertia may misstate the potential for horizontal market power and thereby distort merger enforcement decisions.

These findings highlight the need for empirical models to appropriately represent firms' dynamic pricing incentives and the structure of the post-merger firm. To incorporate these features in settings where firms are asymmetric and face time-series variation in marginal costs and demand, we estimate demand using panel data and then recover the dynamic component of pricing behavior on the supply side. We use this framework to simulate the price effects of horizontal mergers in retail gasoline markets, where we document dynamics consistent with consumer inertia.

On the demand side, our model is a myopic state-dependent extension of logit demand that is close in spirit to Dubé et al. (2009) and Dubé et al. (2010). The central empirical challenge is that state dependence and persistent unobserved heterogeneity can both generate serial correlation in observed choices. With aggregate data, the problem is especially acute because the latent distribution of consumer types is not observed directly. We therefore make a deliberate modeling choice: we restrict static unobserved heterogeneity so that state dependence can be identified from market-level shares. This constrains substitution patterns and can cause the model to overstate the role of inertia relative to a model with richer static preferences. We discuss limitations in more detail in Section 4.1. The benefit is that we obtain a transparent two-step approach that is computationally tractable and well suited to the aggregate data typically used in merger analysis. We view the framework as a disciplined special case for studying dynamic incentives and mergers when inertia is empirically relevant.

On the supply side, we recover the derivative of a firm's continuation value as the residual that rationalizes optimal dynamic pricing under the model. We then conduct counterfactuals using a reduced-form approach to approximate this derivative. Relative to alternative approaches (e.g., Bajari et al., 2007), this has the practical advantages of substantially reducing the computational time of re-computing equilibria and allowing prices to be continuous, rather than projected onto a dimension-reducing grid. Instead of using an approximation to a policy function to recover dynamic and static parameters, we use the structural model to recover demand parameters, and then approximate the dynamic incentives.

As a case study, we apply the model using data from retail gasoline markets, which have advantageous characteristics for our application. The product is relatively homogeneous, transaction costs are minimal, and, to the extent there is inertia, it may be considered relatively short-lived. In terms of dynamics, retail gasoline may represent a conservative test case for our framework: if dynamic incentives matter here, they may be even more important in markets with longer purchase cycles or greater brand differentiation.[4] One limitation is that, by ag-

[4] Airlines are a natural example. Loyalty programs generate state dependence, the industry has experienced

gregating to the market level, our model abstracts from spatial differentiation among stations, which can be important (Houde, 2012).

We present reduced-form evidence of persistence in brand choices in retail gasoline markets, which we interpret as consistent with consumer inertia. Using consumer-level purchase histories, we find that after purchasing from the same brand at least three times in a row, a household purchases from the same brand in the next period 90 percent of the time. However, if they interrupt that spell by buying from a different brand, they return to the previous brand only 57 percent of the time. We also find that consumers are less likely to return to the previous brand if there is a longer period between purchases. These patterns can arise from state-dependent demand, which is the focus of our model, but we recognize that stable, idiosyncratic consumer preferences could generate similar patterns. On the supply side, we show that firms appear to internalize dynamic pricing incentives: new entrants initially set lower prices than established firms and then raise prices over time, and firms begin raising prices in anticipation of predictable future cost increases.

Using our structural model and the restrictions discussed above, we estimate that 63 percent of consumers are subject to state dependence and become affiliated with the brand they purchased most recently.[5] The remaining 37 percent are "shoppers" who are unaffected by consumer inertia. Consumers who are affiliated with a brand are much less price sensitive, with an average elasticity of -0.53. Shoppers are price sensitive, with an average elasticity of -5.96. Across all consumer types, firms face an average elasticity of -1.86.

To compare predictions in merger analysis, we simulate counterfactual mergers between two major gasoline retailers. With the dynamic model, brand consolidation would increase prices for the merging firms by 2.4 percent post-merger. A static model of brand consolidation, by contrast, predicts an average price increase of 5.4 percent, which could attract greater antitrust scrutiny. In this case, the dynamic incentive to invest in future demand mitigates the increase in horizontal market power obtained after a merger. By comparison, the dynamic model also predicts that a merger with joint pricing for separate brands results in a price increase of over 4.9 percent for the merging firms. Thus, modeling the precise structure of the post-merger firm is important for predicting the magnitude of price effects.

Competition authorities often analyze mergers in markets that are likely to be characterized by consumer inertia. For example, in its lawsuit against Swedish Match–National Tobacco, the Federal Trade Commission (FTC) cited strong brand loyalty as a barrier to entry.[6] Similarly, the U.S. Department of Justice cited customer switching as an important factor in its case

substantial consolidation, and longer purchase cycles can increase the lifetime value of retaining a customer. The same aggregate-data logic could therefore be useful in airline merger analysis and related settings.

[5]The National Association of Convenience Stores, a retail fueling lobbying association, found in its 2018 annual survey that 57 percent of respondents have a preference for a specific brand to fill up gasoline. See, https://www.convenience.org/Topics/Fuels/Documents/How-Consumers-React-to-Gas-Prices.pdf

[6]See, "Commentary on the Horizontal Merger Guidelines," 2006, https://www.justice.gov/atr/file/801216/download

against the UPM-MACtac merger. In defense of its acquisition of TaxACT, H&R Block cited the importance of dynamic incentives as a source of downward pricing pressure post-merger (Remer and Warren-Boulton, 2014). Consumer inertia is a central feature of many technology markets currently under heavy scrutiny by antitrust agencies. Yet, perhaps due to computational complexity or compressed investigative timelines, the supply-side implications of inertia are seldom quantified in these investigations.

We consider the implications of consumer state dependence on the pricing behavior of firms, building on a literature that includes Dubé et al. (2009). Our contribution is to examine the effects of competition and the impact of horizontal mergers in such settings. This complements theoretical work on dynamic price competition with habit formation or switching costs (see, e.g., Farrell and Shapiro, 1988; Beggs and Klemperer, 1992; Bergemann and Välimäki, 2006), as well as work on other drivers of dynamic pricing, including experience goods (Bergemann and Välimäki, 1996), network effects (Cabral, 2011), learning-by-doing (Besanko et al., 2018), and search (Stahl, 1989). Our goal is not to distinguish among these mechanisms, but instead to put forward a tractable empirical model that can be used to quantify competitive effects when such dynamics are present.

We contribute to the empirical literature that estimates state dependence in consumer preferences. Meaningful consumer inertia has been documented in packaged goods (Shum, 2004; Dubé et al., 2010; Bronnenberg et al., 2012; Eizenberg and Salvo, 2015), health insurance (Handel, 2013), auto insurance (Honka, 2014), and retail electricity (Hortaçsu et al., 2017); Bronnenberg et al. (2019) provides a summary of the marketing literature on brand loyalty and switching costs. Many of these papers assume that consumers are myopic, which we also maintain. Conceptually, our demand model shares similar features to that of Dubé et al. (2009) and Dubé et al. (2010), though we sacrifice flexibility in static substitution patterns to accommodate market-level price and quantity data. More broadly, empirical models of dynamic demand have addressed stockpiling of storable goods (Hendel and Nevo, 2013) and purchases of durables (Gowrisankaran and Rysman, 2012; Lee, 2013) with forward-looking consumer behavior. The literature highlights that misspecified static models can produce biased elasticities, and Hendel and Nevo (2013) point out that this will matter in a merger analysis. We complement this point by showing that dynamic incentives, rather than biased elasticities, can be the primary concern.

Our identification strategy should be viewed as a tractable special case within a broader literature. Dynamic demand with aggregate data has been studied by Melnikov (2000), Carranza (2010), Gowrisankaran and Rysman (2012), and Lee (2013). In settings with switching costs or state dependence, Ho (2015), Ho et al. (2025), and Shcherbakov (2016) provide more general frameworks. Models that combine state dependence with richer heterogeneity can in principle be identified using additional sources of variation, such as product entry or differentiation moments (Gandhi and Houde, 2019), but not with the exact inversion underlying our two-step procedure. Our objective is to show how a more narrowly focused model can be taken to aggregate data and used for merger analysis.

On the supply side, we propose a reduced-form method to approximate dynamic pricing incentives, allowing us to sidestep some of the challenges present in the estimation of dynamic games. Compared to the value-function approximation methods of Bajari et al. (2007) and Pakes et al. (2007), we lean more heavily on demand-side structure and impose weaker assumptions on firm behavior. Specifically, we evaluate the static demand derivatives using a standard model of demand, and we plug these estimates into the firm's first-order condition to recover dynamic pricing incentives.

The paper proceeds as follows: In Section 2, we introduce the model. Section 3 presents the data for our empirical application and reduced-form evidence of dynamics. Section 4 provides identification arguments and results for dynamic demand estimation. In Section 5, we introduce our approach to evaluate dynamic pricing incentives and conduct a merger analysis. Section 6 concludes.

2 A Model of Oligopolistic Competition with Consumer Inertia

We develop a dynamic model of oligopolistic competition with product differentiation in which consumers may become affiliated with the product they purchased previously. Affiliation may be interpreted as habit formation, brand loyalty, or switching costs.[7] We place parametric restrictions on the form of affiliation for empirical tractability. Consumers in the model are myopic in that they maximize current-period utility rather than a discounted flow of future utility. This assumption is reasonable for retail gasoline markets, as consumers are not likely to choose a gas station anticipating that this will limit their future options. Despite this, some consumers are likely to return to the same gas station due to habit formation, brand loyalty, or inattention.

As detailed below, we introduce consumer dynamics by allowing for a parsimonious form of state dependence in a differentiated product demand model. We then place the demand model into a dynamic oligopoly setting. Even though consumers are myopic, supply-side dynamics arise when firms internalize the effect of sales today on future profits through the accumulation of affiliated consumers.

2.1 Demand

We extend the logit discrete choice model to allow choice probabilities to depend on past purchases.

[7]For certain parameter values, the model can also be interpreted as a model of search or inattention.

Assumption 1: Myopic Discrete Choice Consumers in each market select a single product $j \in J$ that maximizes utility in the current period, or they choose the outside good (indexed by 0). In this setting, we will use "product" and "brand" interchangeably. Each consumer i is indexed by a discrete type, h_i, and a time-varying state, z_{it}. The first feature, h_i, captures exogenous and persistent unobserved heterogeneity.[8] The second feature, z_{it}, allows the distribution of preferences to change endogenously over time through state-dependent utility.

Consumer i receives the following utility for choosing product j in period t:

$$u_{ijt}(z_{it}) = \delta_{jt} + \sigma_{jt}(z_{it}; h_i) + \epsilon_{ijt}. \tag{1}$$

Consumers receive an additively separable common component δ_{jt}, a state-dependent shock that may be type-specific $\sigma_{jt}(z_{it}; h_i)$, and an idiosyncratic shock, ϵ_{ijt}. The common component takes the form $\delta_{jt} = \xi_{jt} + \alpha p_{jt}$ in the standard logit model (with $\alpha < 0$). The mean utility δ_{jt} may depend on time-varying observable characteristics as well as fixed effects. In the empirical application, we make use of this latter feature to help account for serial correlation in unobservable utility shocks over time.

We denote the probability that a consumer of type h in state z chooses product j as $s_{jt}(z; h)$. We normalize $\delta_{0t} = 0$ for the outside good. Given the standard assumption of a type 1 extreme value distribution on the utility shock, ϵ_{ijt}, the choice probabilities are:

$$s_{jt}(z; h) = \frac{\exp(\delta_{jt} + \sigma_{jt}(z; h))}{\sum_{l \in \{0\} \cup J} \exp(\delta_{lt} + \sigma_{lt}(z; h))}. \tag{2}$$

The overall share of product j, denoted S_{jt}, is given by the weighted average of choice probabilities for consumers across states and types.

Assumption 2: Consumer Types Our framework allows for both endogenous and exogenous unobserved heterogeneity. Unobserved exogenous heterogeneity enters our model via two latent consumer types, $h \in \{0, 1\}$. Consumers with $h = 0$ are unaffected by state dependence, so that $\sigma_{jt}(z; 0) = \sigma_{jt}(z'; 0)\, \forall z, z'$, and we normalize $\sigma_{jt}(z; 0)$ to zero for all z. We term these consumers "shoppers." Demand from shoppers is given by the standard logit choice probabilities.[9]

Consumers with a latent type $h = 1$ can be affected by state dependence. The fraction of consumers of this type is given by λ.

[8] The discrete type assumption for the random coefficient model is made elsewhere in the literature. See, for example, Berry et al. (2006) and Berry and Jia (2010).

[9] It is well known that logit demand restricts consumer substitution patterns to be proportional to market share. We impose this restriction to make progress on identifying and estimating a model with state dependence with only market-level data. For applications, such as antitrust investigations, market-level data are common, while more detailed data about diversion are often not available (Valletti and Zenger, 2021). The imposition that diversion is proportional to share may be palatable in more narrow product markets (Miller and Sheu, 2021). In practice, antitrust agencies have employed this assumption during investigations and in court.

Assumption 3: Single-Product Affiliation We now place restrictions on the state-dependent demand shocks for consumers who are affected by state dependence. We assume that each consumer state corresponds to an affiliation (utility shock) to a single product. A consumer in state $z = j$ is affiliated with product j and receives a perceived benefit $\overline{\sigma}_{jt}$, uniform relative to other products: $\overline{\sigma}_{jt} = \sigma_{jt}(j;1) - \sigma_{jt}(z';1)\,\forall z' \neq j$.

We define *affiliation* to be a product-specific state dependence in preferences. The model can be interpreted as brand loyalty when $\overline{\sigma}_{jt}$ is a positive level shock reflecting an internal benefit for purchasing from the same brand, or as a switching cost when $\overline{\sigma}_{jt}$ represents the costs of switching to another brand. These interpretations are empirically indistinguishable because only relative utilities affect choices. In the special case where $\overline{\sigma}_{jt}$ renders affiliated consumers inelastic, the model has a search or inattention interpretation: unaffiliated consumers engage in search while affiliated consumers are inattentive and simply buy the previous product. Distinguishing among these mechanisms lies outside the scope of this paper but may be important for welfare analysis, as brand loyalty and switching costs can produce identical market outcomes but have divergent welfare implications.

We assume that consumers become affiliated with the product they purchased in the previous period. Hence, affiliation is only a function of a consumer's previous choice, rather than a longer purchase history. If a consumer chose the outside option ($j = 0$) in the previous period, the consumer loses their brand affiliation and transitions into the unaffiliated state. This modeling choice differs from approaches in the previous literature (e.g., Keane, 1997), where consumers who do not purchase retain their previous state. We discuss the implications for identification in Section 4.1. In Section 3, we present evidence that our assumption is consistent with purchasing patterns in retail gasoline markets.

Based on these assumptions, we represent the state of each market by the vector $r_t = \{r_{jt}\}$, where r_{jt} denotes the fraction of $h = 1$ consumers affiliated with product j in period t. The share of state-dependent consumers affiliated with product j in period $t+1$ is:

$$r_{j(t+1)} = \sum_{z \in \{0\} \cup J} r_{zt} s_{jt}(z;1). \tag{3}$$

The evolution of states follows a Markov process, where the state can be expressed as a function of the joint distribution of states, types, and choices in the previous period.

Assumption 4: Demand Specification We specify $\delta_{jt} = \xi_{jt} + \alpha p_{jt}$, where ξ_{jt} is composed of product and time (and, later, market) fixed effects. We assume that the affiliation shock affects the utility level by a constant amount for all products, $\overline{\xi}$, while setting the affiliation shock for the outside good to 0. Thus, $\overline{\sigma}_{0t} = 0$ and $\overline{\sigma}_{jt} = \overline{\xi}\,\forall j > 0$. Further, we assume that $\sigma_{jt}(z';1) = 0\,\forall z' \neq j$, so that type $h = 1$ consumers receive the same utility for products they are not affiliated with as shoppers. This rules out arbitrary persistent differences in preferences between the two types of consumers and implies that a consumer of type $h = 1$ who chooses the outside option has the same choice probabilities as a shopper in the subsequent period. Because of this equivalence, we suppress types in the choice probability expression $s_{jt}(z) = s_{jt}(z; h)$ going forward, denoting the choice probabilities of shoppers as $s_{jt}(0)$. We also refer to the combined set of shoppers and consumers with $z = 0$ as *unaffiliated*.

We thus represent the utility of a consumer i in state z for product $j > 0$ as:

$$u_{ijt}(z_{it}) = \xi_{jt} + \alpha p_{jt} + \mathbb{1}[j = z_{it}]h_i\overline{\xi} + \epsilon_{ijt}. \tag{4}$$

We represent $\overline{\sigma}_{jt} = \mathbb{1}[j = z_{it}]h_i\overline{\xi}$, where consumers subject to affiliation have $h_i = 1$ and shoppers have $h_i = 0$. The error term ϵ_{ijt} follows the type-1 extreme value distribution, yielding the choice probabilities specified in equations (2) and (3).

The aggregate share of product j across consumer types and states is therefore:

$$S_{jt} = (1 - \lambda)s_{jt}(0) + \lambda \sum_{z=0}^{J} r_{zt}s_{jt}(z). \tag{5}$$

The aggregate share can thus be written as a weighted sum of product j's share of unaffiliated consumers, $s_{jt}(0)$, and affiliated consumers, $s_{jt}(z)\forall z \neq 0$. Note that the share of unaffiliated consumers in any period is $(1 - \lambda) + \lambda r_{0t}$, as some fraction of state-dependent consumers may have chosen the outside option in the prior period.

2.2 Supply

We assume that firms set prices to maximize the net present value of profits. We restrict attention to Markov perfect equilibria.

Assumption 5: Competition in Prices Firms set prices in each period to maximize the net present value of profits from an infinite-period game. Prices are set as a best response conditional on the state and contemporaneous prices of rival products. Firms cannot commit to future prices. The state vector is summarized by marginal costs, c_t, the distribution of affiliation across consumers, r_t, and other variables captured by x_t. Entry is exogenous. The objective function for firm k can be summarized by the Bellman equation:

$$V_k(c_t, r_t, x_t) = \max_{p_{kt}|p_{-kt}} \big\{(p_{kt} - c_{kt})S_{kt} + \beta E(V_k(c_{t+1}, r_{t+1}, x_{t+1})|p_t, c_t, r_t, x_t)\big\}. \tag{6}$$

Prices optimize the sum of current-period profits $(p_{kt} - c_{kt})S_{kt}$ and the continuation value. When the perceived continuation value is non-zero, firms anticipate how price affects the future distribution of consumer states and the impact of future changes to marginal costs. Note that the state space does not include previous-period prices, excluding strategies that depend directly upon competitors' histories.

Assumption 6: Expectations Consistent with the Markov perfect framework, we assume that the continuation value function is stable conditional on the state and prices. In contrast to the typical setup for a dynamic game, we place minimal restrictions on expectations, discount rates, and the perceived continuation value. Instead, our empirical approach, which we describe in Section 5, directly estimates a reduced-form model of the derivative of the continuation value.

Market equilibrium is characterized by consumers making (myopic) utility-maximizing purchase decisions and firms pricing as the best response to other firms' prices, conditional on the state.

Merger Implementation

We consider two types of horizontal mergers. The first, a *joint pricing* merger, unites pricing control of two products under a single firm while maintaining them as separate products and brands. The second type, *brand consolidation*, consolidates two products under a single brand, effectively offering a single product post-merger.

In a joint pricing merger, we assume that the merged firm maximizes the objective function in equation (6) while setting prices for separate brands. In the first post-merger period, we assume consumers remain affiliated with the brand that they were affiliated with in the last pre-merger period. The post-merger firm's objective function changes because it internalizes how the prices of each of its two products affect the profits of both.

To implement a brand consolidation merger, we assume that consumers affiliated with the removed brand transition to state 0 following the merger. We also assume that, at pre-merger prices, the consolidated brand will have the same combined share of shoppers as the separate pre-merger brands. We implement this by adjusting the value of ξ_j for the remaining brand. This assumption facilitates an "apples-to-apples" comparison to joint pricing mergers. In the context of retail gasoline markets, this is akin to assuming that shoppers derive value from features such as location rather than from the brand per se. One implication is that at pre-merger prices, the share of affiliated customers increases for the remaining brand relative to the two separate brands, giving the merged firm an advantage in retaining affiliated consumers.

The presence of demand dynamics is a key factor that leads to differential outcomes across these two merger types. By contrast, in a symmetric oligopoly with static logit demand, these two merger types yield identical price effects. See Appendix A for a proof of this result and for additional technical details about implementation.

3 Data and Reduced-Form Evidence of Dynamics

We now introduce the data used in the demand estimation and empirical application, and provide evidence of dynamic demand and dynamically adjusting retail gasoline prices. Appendix B reports additional reduced-form evidence on demand persistence and pass-through dynamics.

3.1 Data

The main empirical analysis relies on daily retail prices for regular fuel at nearly every gas station in Kentucky and Virginia, totaling almost six thousand stations. As a measure of marginal cost, the data include the brand-specific daily wholesale rack price charged to each retailer, as well as federal, state, and local taxes. We therefore almost perfectly observe each gas station's marginal-cost changes, except for privately negotiated per-gallon discounts that are likely fixed over a year. The data span October 2013 through September 2015. These data were obtained from the Oil Price Information Service (OPIS), which has previously provided data for academic studies (e.g., Lewis and Noel, 2011; Chandra and Tappata, 2011; Remer, 2015).

OPIS also supplied market share data, reported by week and county for each gasoline brand.[10] In our analysis, we treat consumers as choosing among brands in a county. Due to contractual limitations, OPIS only provided inside market shares, not actual volume. To account for temporal changes in market-level demand, we supplement the share data with monthly, state-level consumption data from the Energy Information Administration (EIA). We describe our adjustment in Section 4.3. To account for differences in demographics across markets, we merge the data with measures of income and population density from the American Community Survey.

We use an additional NielsenIQ dataset to document retail gasoline purchasing patterns and provide suggestive evidence of consumer inertia. We employ the Consumer Panel Data to track individual household purchasing decisions over time. These data do not include prices or provide comprehensive market estimates. We use these data only to document patterns of repeat purchases.

3.2 Dynamic Demand: Evidence from Consumer Data

We use the NielsenIQ Consumer Panel Data to analyze household purchases of gasoline from 2007 through 2018. We exclude very small and very large purchase amounts (under \$10 or over \$120), exclude gasoline purchases from warehouse clubs and grocery stores, and limit the analysis to households with at least 26 purchases per year. Within this sample, the median time between purchases is 7 days, and in over 90 percent of cases, the household purchases only

[10] OPIS calculates these data from actual purchases that fleet drivers charge to company cards issued by Wright Express (WEX), the largest provider of fleet cards in the United States with over 4 million drivers and 95 percent coverage of fuel retailers.

Figure 1: Fraction of Repeat Purchases by Household

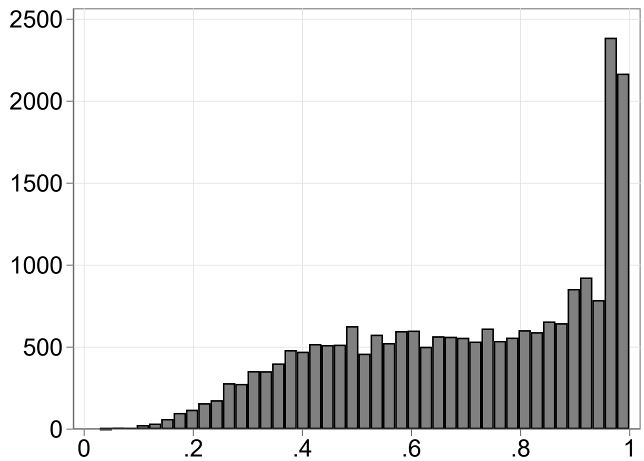

Notes: This figure depicts the distribution of households by fraction of repeat purchases. For each household we calculate the average fraction of gasoline purchases that are a repeat purchase, defined as returning to the same brand as the previous purchase. We restrict the sample to households with at least 26 gasoline purchases in a year.

gasoline during the trip. In our sample, 63 percent of trips are from households that meet our threshold restrictions; we obtain qualitatively similar findings without these restrictions.

Using these data, we explore the propensities of households to return to the same brand. For each household, we calculate the fraction of gasoline purchases that are repeat purchases (i.e., the same brand as the previous purchase). Figure 1 plots this distribution. Three features indicate persistence in brand choices. First, repeat purchases are exceedingly common: for the median household, 73 percent of trips are repeat purchases. Second, there is a mass of households that shop almost exclusively at a single brand: the 90th percentile household has 98 percent repeat purchases. Third, there is substantial variation in the tendency to make repeat purchases. This pattern is consistent with different "types" of households; some may exhibit strong state dependence or brand affiliation while others shop around. We interpret these features as suggestive evidence of consumer inertia, as static heterogeneity in brand preferences could yield similar patterns. Appendix B.2 complements this household-level evidence with brand-level regressions showing that lagged shares remain strongly predictive of current shares even after rich fixed effects.

Next, we use the panel data to examine household-level choices over time. We identify all spells where a household purchases from the same brand for at least 3 consecutive trips (156,150 distinct spells). In these spells, the probability that the next trip is a repeat purchase is 0.903. After the spell ends and a different brand has been chosen, the probability of returning to the prior brand is 0.565—much lower than the probability of staying conditional on purchasing in the previous period. The probability of staying with the new brand is 0.266,

Table 1: Repeat-Purchase Patterns in Shopping Behavior

Variable	Value
Probability of Repeat Purchase After At Least Three Consecutive Trips	0.903
Conditional on Switch, Probability of Switching Back to Prior Brand	0.565
Conditional on Switch, Probability of Staying with New Brand	0.266
Conditional on Switch, Combined Probability of Staying with New Brand or Switching Back	0.831

Notes: Each statistic is calculated from spells that have at least three consecutive trips to the same brand.

while the probability of purchasing from any other brand is 0.169. These patterns, which are summarized in Table 1, are consistent with state-dependent demand where utility depends on the brand chosen in the previous period. We note that this evidence is not a direct test of state dependence—persistent, idiosyncratic preferences could also generate these patterns, with households switching due to, e.g., price changes.

We also leverage data on the duration between purchases. Consistent with our model where consumers who choose the outside option lose their affiliation, households are less likely to return to the same brand after a longer interval. This effect is strongest for households that make repeat purchases frequently, and is zero for those least likely to make repeat purchases—consistent with these latter consumers being "shoppers" that are unaffected by inertia. Further details are provided in Appendix B.1.

3.3 Dynamic Pricing

We now present reduced-form evidence of dynamic pricing by retail gasoline stations. Consistent with a model where firms accumulate affiliated consumers over time, we find that new entrants price lower than established competitors, and this discount dissipates over time. We also examine cost pass-through, showing that firms are slow to adjust to marginal cost changes and anticipate expected cost increases by raising prices in advance.

3.3.1 Dynamic Pricing of New Entrants

When forward-looking firms price to consumers who may become affiliated, there is an incentive to initially offer prices below the static optimum in order to build up affiliated customers. Using the station-level OPIS data, we identify 193 new entrants.[11] We compare the prices of new entrants to the county-average price of incumbent stations.

Figure 2 shows that gas stations enter with prices on average two cents per gallon lower than incumbents, then slowly converge to the market average. For the first 8 weeks following entry, new entrant prices are 2.1 cents per gallon lower than incumbents' (standard error:

[11] To ensure there are sufficient data and to control for composition effects, we limit the set of entrants to stations with at least one year of post-entry price data.

Figure 2: New Entrant Prices

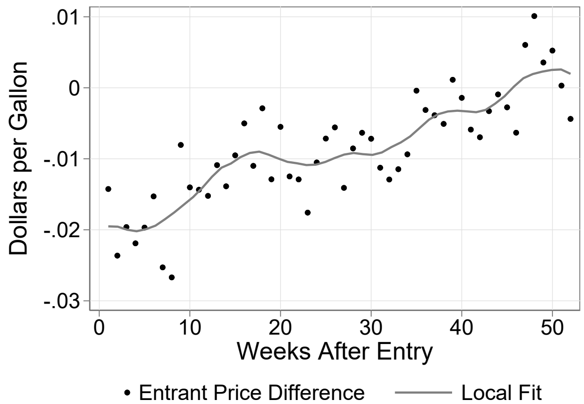

Notes: A data point measures the average difference between a new entrant's price and the county average price, for the given number of weeks after entry. The line is created using local polynomial regression.

0.24). From weeks 9 through 24, entrant prices are 1.1 cents per gallon lower (standard error: 0.17). These differences are highly significant and, given our demand estimates, economically meaningful for attracting unaffiliated consumers.

3.3.2 Cost Pass-through

When pricing to consumers prone to inertia, a profit-maximizing firm will slowly adjust price in response to a cost change. Similarly, it will begin increasing price in expectation of a future cost change. Smoothing a price change over time is optimal, as a large immediate price increase would drive away too many affiliated consumers.

To highlight the temporal component of cost pass-through, we separately estimate how gas stations react to expected versus unexpected cost changes. We construct expected cost using gasoline futures and current wholesale costs to project 30-day-ahead costs; unexpected costs represent deviations from this projection.[12] We estimate the following model:

$$p_{nt} = \sum_{s=-50}^{50} \beta_s \hat{c}_{n(t-s)} + \sum_{s=-50}^{50} \gamma_s \tilde{c}_{n(t-s)} + \sum_{s=-50}^{50} \phi_s \tau_{n(t-s)} + \psi_n + \varepsilon_{nt}. \quad (7)$$

Here, p_{nt} is the price at gas station n at time t, $\hat{c}_{n(t-s)}$ and $\tilde{c}_{n(t-s)}$ are expected and unexpected wholesale costs with lag s, and $\tau_{n(t-s)}$ is the state-level sales tax. We construct cumulative response functions to track price adjustment to a one-time, one-unit cost change at time $t = 0$.

Figure 3 plots the cumulative response functions. For unexpected costs (Panel a), prices react suddenly at time zero but take about four weeks to reach the full response. Estimated

[12] For details, see Section B.3 in the Appendix.

Figure 3: Cumulative Pass-through

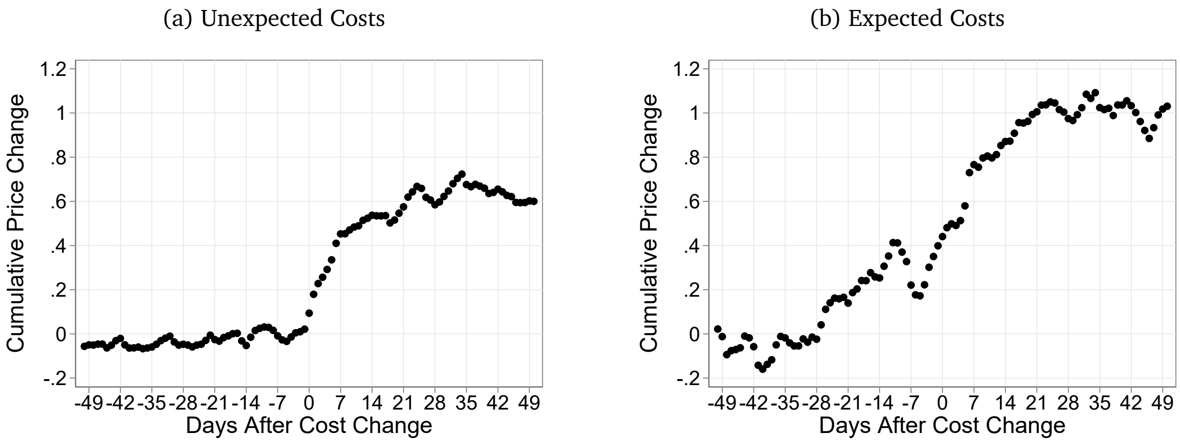

Notes: Panels (a) and (b) depict the cumulative price change in response to a one-unit cost change at time = 0. Response functions are created from the estimated parameters of equation (7).

pass-through peaks at 0.72 after 34 days, averaging 0.64 over days 21 through 50. For expected costs (Panel b), firms begin reacting approximately 28 days in advance, with prices capturing about 40 percent of the effect the day before the cost arrives. The estimated pass-through rate averages 1.01 over days 21 through 50. Long-run pass-through rates differ across the type of shock: expected costs experience approximately "full" pass-through, while unexpected costs demonstrate incomplete pass-through of about 64 cents per dollar.

These patterns—slow adjustment to marginal cost changes and anticipation of expected changes—are consistent with forward-looking behavior by firms and dynamic demand arising from consumer affiliation.[13]

4 Demand Analysis: Identification and Estimation

Given the reduced-form evidence of dynamic demand and supply behavior, we now present the empirical application of the model to retail gasoline markets. We divide estimation into two stages, as demand can be estimated independently of supply-side assumptions. Our method of demand estimation relies on data widely used in static demand estimation: shares, prices, and an instrument. In Section 5, we use the estimated demand system to analyze the dynamic incentives faced by suppliers.

[13] In robustness checks, we find little evidence of price asymmetry between positive and negative cost shocks, and no evidence of Edgeworth price cycles.

4.1 Identification

We organize the identification argument in three steps. First, conditional on the parameters governing state dependence, the structure of the model lets us recover the latent choice probabilities $\{s_{jt}(z)\}$ from observed shares. Second, those latent shares allow us to recover mean utility for unaffiliated consumers and estimate the static demand parameters using standard instrumental variables (IV) arguments. Third, serial-correlation moments identify the dynamic parameters.

This strategy should be interpreted as a particular aggregate-data implementation of the broader literature on dynamic demand and switching costs, which includes Ho (2015), Ho et al. (2025), and Shcherbakov (2016). By restricting the form of static unobserved heterogeneity, we obtain exact recovery of latent choice probabilities from observed shares. The payoff from that restriction is a two-step estimator that is transparent, computationally tractable, and feasible with the market-level data commonly used in merger analysis.

A key challenge with aggregate data and unobserved heterogeneity is that we do not separately observe choice patterns by unobserved consumer type. In our context, we observe the aggregate share, S_{jt}, which is a weighted combination of the type-specific choices $\{s_{jt}(z)\}$ and depends on the distribution of affiliated consumers for each product $\{r_{jt}\}$. Observed aggregate shares are given by (5).

To separate out $\{s_{jt}(z)\}$ from S_{jt}, we leverage the structure of the model. With discrete types, we show exact identification of the choice distribution without supplemental assumptions.

Proposition 1 *With discrete types, the distribution of choice patterns is identified conditional on the distribution of types and type-specific shocks.*

Using the dynamic extension of the logit demand system detailed in Section 2, we obtain the familiar expression for the log ratio of shares of unaffiliated consumers from equation (2):

$$\ln s_{jt}(0) - \ln s_{0t}(0) = \delta_{jt} \tag{8}$$

Likewise, we obtain the following relation for shares of affiliated consumers:

$$\ln s_{jt}(z) - \ln s_{0t}(z) = \delta_{jt} + \sigma_{jt}(z). \tag{9}$$

To show identification, we substitute equation (8) into (9) and use the fact that $\frac{1}{s_{0t}(j)} - \frac{1}{s_{0t}(0)} =$

$\exp(\delta_{jt})(\exp(\sigma_{jt}(j)) - 1)$ to obtain:

$$s_{jt}(z) = \exp(\sigma_{jt}(z)) \frac{s_{0t}(z)}{s_{0t}(0)} s_{jt}(0) \tag{10}$$

$$s_{jt}(0) = \left(\frac{s_{0t}(0)}{s_{0t}(j)} - 1\right) \frac{1}{\exp(\sigma_{jt}(j)) - 1}. \tag{11}$$

Thus, the $J + J^2$ unknowns $\{s_{jt}(z)\}_{|j \neq 0}$ can be expressed in terms of the $J + 1$ unknowns $\{s_{0t}(j)\}$ and $s_{0t}(0)$. These $J + 1$ unknowns are pinned down by the adding-up condition $1 - \sum_j s_{jt}(0) - s_{0t}(0) = 0$ and the observed share equations given by (5).

Identification requires $\{r_{jt}\}$, which is the state describing the share of state-dependent consumers affiliated with each product. Given our assumptions about the evolution of demand, the value of r_{jt} can be calculated from prior-period values of $S_{j(t-1)}$ and $s_{j(t-1)}(0)$:

$$r_{jt} = \frac{1}{\lambda}\left(S_{j(t-1)} - (1 - \lambda)s_{j(t-1)}(0)\right), \tag{12}$$

so, given an initial $\{r^*_{j0}\}$, the state can be iterated forward market by market and period by period. Thus, conditional on the dynamic parameters and the initial state, the data identify the latent shares relevant for the static demand step.

We now incorporate multiple markets in our notation, which we index by m. From (8), the inversion recovers the mean utility of the unaffiliated consumer in each market, which we specify as a linear function of covariates:

$$\delta_{jmt} = X_{jmt}\gamma + \eta_{jmt}. \tag{13}$$

This is analogous to recovering mean utility in Berry et al. (1995). The covariates X_{jmt} may incorporate endogenous variables, such as price, and are identified by standard IV arguments. This is a significant practical benefit of the two-step procedure: once the latent shares of unaffiliated consumers are recovered, the inner problem is a linear IV regression that can be estimated quickly and transparently.

To identify λ_m and $\sigma_{jt}(j)$, we employ additional moments. We use fixed effects in X_{jmt} to model unobserved serial correlation in demand and calculate the residual demand innovations η_{jmt} after accounting for these fixed effects. Specifically, in our empirical specification, we include fixed effects that capture aggregate period-specific demand shocks, product-specific persistent demand, and market-specific seasonal patterns. We then assume the residual demand innovations η_{jmt} are serially uncorrelated, so that $E[\eta_{jmt}\eta_{jm(t+1)}] = 0$ holds on average within each brand.

We parameterize $\lambda_m = f(D_m\theta)$. Thus, the share of consumers subject to state dependence can vary with market-level demographic measures D_m. We previously restricted $\sigma_{jt}(j)$ to equal $\overline{\xi}$ for all $j \neq 0$ and 0 for $j = 0$, such that affiliated customers receive a constant level shock to utility (Section 2.1). Separate identification of θ and $\overline{\xi}$ is made possible by the structure of the model. The share of consumers who become affiliated, λ_m, does not depend on price, whereas the impact of $\overline{\xi}$ on shares does. As can be seen by examining equations (8) and (9), a change in price affects δ_{jmt}, which shifts relative choice patterns holding fixed $\overline{\xi}$. Each component of θ is identified by how serial correlation patterns covary with market-level demographics. The presence of entry aids identification: new brands have zero affiliated consumers when they enter the market, and the estimated parameters reflect how entrants' prices and shares evolve as they accumulate affiliated consumers. In the 241 markets in our data, 90 experience entry of a new brand. Appendix C.2 provides Monte Carlo evidence that, under the maintained structure, the estimator can distinguish state dependence from persistent heterogeneity.

Limits of Approach

The identifying assumption in our model is not that inertia is the only possible source of persistence in observed shares. Rather, it is that after controlling for fixed effects and common shocks, the remaining serial structure relevant for brand-market demand can be attributed to the low-dimensional affiliation process in the model. If there are brand-market-specific transitory shocks with their own persistence, then some of that persistence may load onto the estimated dynamic parameters. In that sense, our estimates should be interpreted as the component of demand persistence that the affiliation model can rationalize given aggregate price and share data.

If non-shopper consumers also had idiosyncratic and persistent brand-specific tastes, the log-ratio simplification would no longer apply. Such an approach would require additional types $h \in \{1, \ldots, H\}$ that yield aggregate shares $S_{jt} = (1-\lambda)s_{jt}(0) + \lambda \sum_{h=1}^{H} \sum_{z=0}^{J} r_{hzt} s_{hjt}(z)$. The average choice probability across types within a state would be $\overline{s}_{jt}(z) = \frac{1}{\overline{r}_{zt}} \sum_{h=1}^{H} r_{hzt} s_{hjt}(z)$, where $\overline{r}_{zt} \equiv \sum_{h=1}^{H} r_{hzt}$, and the expression $\ln \overline{s}_{jt}(z) - \ln \overline{s}_{0t}(z)$ is a complicated function of δ_{jt} and $\{\sigma_{jt}(z;h)\}$ that cannot be simplified analytically.

This is where our approach differs from richer frameworks such as Ho (2015), Ho et al. (2025), and Shcherbakov (2016). More general models can in principle combine state dependence and static heterogeneity, but they require additional sources of variation—such as product entry and exit, excluded-variable dynamics, observed heterogeneity, or Gandhi and Houde (2019) differentiation-style moments—rather than the exact inversion that underlies our two-step procedure. We therefore interpret our method as a transparent special case designed for aggregate data and merger work, not as a claim that richer models are unidentified. As a consequence of our identifying restrictions, the model may well overstate the importance of inertia relative to approaches that allow for richer heterogeneity in preferences.

The treatment of the outside option also matters for identification and interpretation. The scanner-data applications in Dubé et al. (2009) and Dubé et al. (2010) focus on persistence across inside-good purchases, Ho (2015) studies attachment to inside goods, and the "start-up" cost assumption of Shcherbakov (2016) also generates inertia for non-purchase. These various assumptions will lead patterns in the data to load on different parameters. In our setting, we conceptualize affiliated consumers as regular drivers who stop at a gas station once each week. When consumers elect the outside option, they "reset" their decision-making and act as if they are unaffiliated. Consistent with this, our reduced-form results show that, within a household, a longer purchase interval is associated with a lower probability of a repeat purchase (see Appendix B.1). If affiliation persisted through outside-option choices, we would not expect to see this relationship. Allowing for two types of consumers—those with a tendency to become affiliated and shoppers who do not have this tendency—gives us additional flexibility. Consumers who do not regularly purchase gasoline each week are captured in our model as shoppers.[14] Thus, the magnitude of the outside-option share and its variation over time help discipline both the mass of shoppers and the switching probabilities of affiliated consumers.

A potential concern with our approach is that if consumers indeed retain strong brand attachment even after long spells of non-purchase, then our model will understate persistence. Likewise, the one-period-history assumption will be too short if affiliation lasts for several purchase cycles. Relaxing that assumption is conceptually straightforward—one can enlarge the state to include longer histories or impose a parametric decay structure—but doing so would substantially expand the latent state space. In practice, that would likely require either consumer-level panel data or stronger exclusion restrictions than we use here.

This discussion also helps frame external validity. Gasoline is attractive because purchases are frequent and the previous purchase plausibly contains a large share of the relevant state. In industries with longer purchase cycles, such as airlines or consumer banking, one might want the state to depend on a longer history, on accumulated status, or on prior product holdings. The same identification logic can still be relevant, but the state representation would likely need to be richer. When applying our approach to other settings, it is important that the time period of aggregation corresponds to a reasonable purchase interval, or that the model is adapted properly to allow for the dynamic considerations to last across multiple periods.

4.2 Estimation Routine

Our estimation routine recovers the dynamic parameters $(\theta, \bar{\xi})$ and the static demand parameters γ. The estimator is a method-of-moments procedure based on the restriction that residual demand shocks are serially uncorrelated after fixed effects are removed. We use both covariance and correlation versions of this restriction to improve robustness to extreme values. Aggregating each set of moments by brand yields $16 \times 2 = 32$ moments in the objective function.[15] We estimate three parameters in θ, yielding four total dynamic parameters to be estimated.

[14] We estimate that the fraction of active consumers each week that are shoppers is 37 percent. To calculate the number of underlying consumers who are shoppers, one would have to scale this fraction by the inverse of the rate of purchases, i.e., doubling it if shoppers only purchase once every two weeks.

[15] As detailed below, we have 16 brands in the data, including a synthetic "fringe" brand.

Formally, the estimate of the dynamic parameters $\tilde{\theta} \equiv (\theta, \bar{\xi})$ is given by

$$\hat{\theta} = \arg\min_{\tilde{\theta}} g(\tilde{\theta})' W g(\tilde{\theta}), \qquad g(\tilde{\theta}) = \begin{bmatrix} g^{cov}(\tilde{\theta}) \\ g^{corr}(\tilde{\theta}) \end{bmatrix}. \tag{14}$$

Here, $g^{cov}(\tilde{\theta})$ and $g^{corr}(\tilde{\theta})$ are the $J \times 1$ sample analogs of the covariance and correlation moments. Additional implementation details are collected in Appendix C.

For a candidate $\tilde{\theta}$, the routine proceeds as follows. First, compute λ_m and $\sigma_{jmt}(j)$. Second, recursively solve the market-period nonlinear system from Section 4.1 conditional on r_{jmt}, starting from the initial value r^*_{jm0}. Third, recover $s_{jmt}(0)$ and construct $\delta_{jmt} = \ln s_{jmt}(0) - \ln s_{0mt}(0)$. Fourth, estimate (13) by IV, obtain the residuals $\hat{\eta}_{jmt}$, and evaluate the moment vector. Fifth, search over $\tilde{\theta}$ to minimize (14).

Two computational simplifications keep the procedure feasible. First, Appendix C.1 shows how the market-period system can be reduced to two unknowns, $s_{0t}(0)$ and $\sum_{z=0}^{J} r_{zt} s_{0t}(z)$, with explicit formulas for the remaining values. Second, the linear form of the inner IV regression makes the nested optimization fast. Because the panel spans 104 periods, the effect of the initial condition is largely absorbed during a burn-in. In the baseline, we set r^*_{jm0} equal to the mean observed share in the period prior to our sample, and robustness checks indicate that we obtain very similar point estimates with alternative starting values.

4.3 Empirical Specification

To construct market shares that allow for the outside option ($j = 0$), we merge the market share data provided by OPIS with monthly state-level consumption data provided by the EIA. We assume that the maximum observed quantity in the EIA data reflects 75 percent of the total potential market, and, for each state-week, we construct a quantity-based scaling factor equal to $0.75 \times q_{st} / \max_t\{q_{st}\}$, where q_{st} is the state-level quantity. We smooth this scaling factor using a centered three-week moving average before applying it to the OPIS inside shares. The outside-option share is the residual needed for shares to sum to one.

We merge these EIA-adjusted brand-county shares with the average prices for each brand in a week-county. To reduce the occurrence of zero shares, we linearly interpolate gaps of up to two weeks and treat longer gaps as periods in which the station is not in the choice set.[16] We drop observations with missing prices, missing shares, or missing lagged shares.

Table 2 provides summary statistics for the 241 counties in Kentucky and Virginia. To reduce sensitivity to small brands and keep the counterfactual exercises tractable, we aggregate selected brand-market observations into a synthetic "fringe" product. A brand is retained as a standalone brand (separate from the fringe) if it appears in at least ten markets and either

[16] After all cleaning steps, approximately 1 percent of observations have zero shares, and over 90 percent of these zero-share occurrences are in spells of six weeks or longer.

Table 2: Summary Statistics by County

Statistic	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max	N
Num. Brands	4.52	1.46	1.00	3.00	6.00	8.00	241
Price	2.87	0.11	2.53	2.77	2.96	3.14	241
Wholesale Price	2.25	0.06	2.02	2.21	2.28	2.44	241
Margin ($/gal)	0.21	0.06	0.06	0.17	0.24	0.44	241
Num. Stations	22.33	27.88	1.88	7.69	25.59	239.13	241

Notes: Table displays summary statistics averaged across each of the 241 markets in the sample. Prices and margins are reported in dollars per gallon.

has an average share of at least 2 percent across all markets or an average share of at least 10 percent for the markets in which it is present. We then impose a market-level screen: if a retained brand's average share within a given market is below 5 percent, we assign that brand-market to the fringe category. We also assign unbranded observations to the fringe category. This reduces the maximum number of brands per county from 24 to 8. Across all markets, we analyze pricing for 16 brands, including the synthetic fringe.[17]

We take steps to reduce measurement error in the number of stations. We assume that stations exist for any gaps in our station-level data lasting fewer than 12 weeks. Likewise, we trim for entry and exit by looking for 8 consecutive weeks (or more) of no data at the beginning or end of the sample. After cleaning, the sample contains 110,844 observations. Summary statistics at the observation level are reported in Appendix Table D.1.

We implement regression equation (13) as:

$$\delta_{jmt} = \alpha p_{jmt} + \pi \left(p_{jmt} \times Income_m\right) + \kappa N_{jmt} + \zeta_{jm} + \phi_t + \psi_{m,month(t)} + \eta_{jmt}. \quad (15)$$

We obtain $\delta_{jmt} = \ln\left(\frac{s_{jmt}(0)}{s_{0mt}(0)}\right)$ following the first three steps of our estimation routine in Section 4.2. The brand-county fixed effects ζ_{jm} absorb time-invariant local demand shifters such as amenities and demographics. We also include the number of stations, N_{jmt}, which is identified by within-brand-county variation over time.

The weekly fixed effects $\{\phi_t\}$ capture aggregate demand shocks, while the county-specific monthly effects $\{\psi_{m,month(t)}\}$ absorb local seasonality. After partialling out these components, the identifying restriction for the dynamic parameters is that the residual demand innovation η_{jmt} is uncorrelated over time.

We allow for price endogeneity by instrumenting for p_{jmt} with variation in wholesale costs driven by U.S. crude-oil production. The first instrument (Z_1) is constructed from a regression of deviations of wholesale costs from the brand-county average on the interaction of weekly U.S. crude-oil production with the brand-county average wholesale cost. We interact this instrument with $Income_m$ to create a second instrument, Z_2, for $(p_{jmt} \times Income_m)$. Appendix Figure D.1

[17] Summary statistics by brand are presented in Appendix Table D.2. The fringe brand accounts for, on average, 13 percent share across markets where it appears.

reports the time series of inside shares, prices, and the first instrument.

For the dynamic parameters, we estimate $\bar{\xi}$ directly and specify $\lambda_m = f(D_m\theta)$ as:

$$\lambda_m = \frac{\exp(\theta_1 + \theta_2 Income_m + \theta_3 Density_m)}{1 + \exp(\theta_1 + \theta_2 Income_m + \theta_3 Density_m)} \tag{16}$$

Thus, we allow the share of consumers subject to state dependence to vary with market-level measures of median household income and (log) population density.[18]

4.4 Demand Estimation Results

The estimates for the linear demand parameters are reported in Table 3.[19] The first three columns report coefficient estimates from a logit demand regression using observed shares, whereas the fourth column reports the results for unaffiliated consumers from our dynamic model. With the full set of fixed effects, the price coefficient is similar in the static and dynamic specifications, though the interpretation differs because the dynamic regression uses latent shares of unaffiliated consumers only. We estimate essentially no relationship between income and price sensitivity for unaffiliated consumers. We also find that an increase in the number of stations has a statistically significant positive effect on demand.

Table 4 reports estimates of the dynamic parameters. The coefficients $(\theta_1, \theta_2, \theta_3)$ imply that, on average, 62.8 percent of consumers are subject to state dependence and develop an affiliation to the brand they most recently purchased from, i.e., $E[\lambda_m] = 0.628$. The coefficient of -0.741 on income indicates that lower-income markets tend to have more consumers prone to affiliation, though this relationship is not statistically significant at the 95 percent level. The statistically significant coefficient of 0.287 on population density indicates a higher share of state-dependent consumers in denser areas.

The estimated utility shock $\bar{\xi}$ implies that affiliated consumers are relatively inelastic with respect to price. Across observations, the mean own-price elasticity for affiliated consumers is -0.53, and the median is -0.38. On average, an affiliated consumer repurchases from their preferred brand 92 percent of the time. Unaffiliated consumers are much more price sensitive, with an average own-price elasticity of -5.96. Thus, a 1 percent increase in price (roughly 3 cents) reduces unaffiliated demand by about 6 percent.

On average, roughly 76 percent of a brand's customers are affiliated consumers in any week in equilibrium. The average weighted elasticity, which combines affiliated and unaffiliated consumers using their equilibrium shares, is -1.86. This weighted elasticity is the relevant elasticity faced by a firm and is substantially less elastic than the "naive" static estimate of -5.74. At the market level, the estimates imply an aggregate weekly elasticity of demand of

[18] Several studies find that consumer behavior in retail gasoline markets is affected by income. See, for example, Nishida and Remer (2015), Levin et al. (2017), and Donna (2021).

[19] At the solution, our parameter estimates deliver an objective close to zero. The implied overall autocorrelation in shocks is -0.0007, and the overall covariance is -0.0003.

Table 3: Estimates of Static Demand Parameters

	Static Model			Dynamic Model
	(1)	(2)	(3)	(4)
Price	-0.022^{*}	-0.260^{***}	-2.315^{***}	-2.198^{***}
	(0.014)	(0.038)	(0.507)	(0.458)
Price $\times$ Income	-0.135^{***}	0.077^{***}	-0.007	0.006
	(0.021)	(0.023)	(0.032)	(0.033)
Number of Stations	0.016^{***}	0.064^{***}	0.063^{***}	0.058^{***}
	(0.004)	(0.010)	(0.010)	(0.010)
IV	No	No	Yes	Yes
Brand-County FEs		X	X	X
Week FEs		X	X	X
County-(Month of Year) FEs		X	X	X
Observations	110,844	110,844	110,844	110,844

Notes: Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Table displays the estimated coefficients for a logit demand system. For the first three models, the dependent variable uses observed aggregate shares. For the fourth model, the dependent variable uses the shares of unaffiliated consumers in the dynamic model. The Sanderson-Windmeijer conditional F-statistics for the excluded instruments are 56.9 for price and 141.2 for price $\times$ income. Standard errors are clustered at the county level. For the dynamic model, standard errors are calculated via the bootstrap.

-1.20, close to the short-run elasticity of -1.38 in Levin et al. (2017), who estimate gasoline demand using daily city-level expenditure and price data.

We assess robustness along two dimensions. First, we vary the initial value of the unobserved state. We consider two extreme alternatives: no initial affiliation ($r^*_{jm0} = 0$) and an upper-bound case in which all non-shoppers are initially affiliated. Both deliver point estimates similar to the baseline, and all dynamic and static parameters remain within the 95 percent confidence intervals of our preferred specification. Second, we test for numerical uniqueness by repeatedly drawing random initial parameter values. The estimator converges to the same solution in each case, up to negligible rounding error.

5 Supply-Side Analysis

In this section, we use our demand estimates in conjunction with the supply model to analyze the impact of consumer inertia on market power and horizontal mergers in an empirical setting. We use the demand estimates from Section 4 together with the supply model detailed in Section 2.2 to simulate the effects of mergers in retail gasoline markets. We study both joint pricing and brand consolidation mergers, demonstrating that this modeling choice has important implications for post-merger outcomes. For additional insight, we then characterize the steady-state

Table 4: Estimates of Dynamic Demand Parameters

	Affiliation Rate			Strength of Affiliation
	Baseline θ_1	*Income* θ_2	*Density* θ_3	*Utility Shock* $\bar{\xi}$
Coefficient	0.584	-0.741	0.287	5.833
95 Percent CI	[0.48, 0.75]	[-1.35, 0.27]	[0.09, 0.51]	[5.00, 6.53]

Notes: Table displays the estimated non-linear coefficients from the dynamic model. The first three parameters imply that, on average, 62.8 percent of consumers who purchase are subject to state dependence. Confidence intervals are shown with brackets and are calculated via the bootstrap.

equilibrium and present numerical simulations. Finally, we discuss antitrust implications of the interaction of consumer inertia and horizontal mergers.

5.1 Quantifying Dynamic Pricing Behavior

Given our demand estimates, we construct the components of each firm's Bellman equation from equation (6). We assume that the price for each brand is set to maximize discounted profits at the county level. Following equation (6), the dynamic condition for optimal pricing for a single-product firm that owns brand j is:

$$\frac{\partial \pi_{jt}}{\partial p_{jt}} + \beta \frac{\partial E\left[V_j(r_{t+1}, c_{t+1}, x_{t+1}) | p_t, r_t, c_t, x_t\right]}{\partial p_{jt}} = 0, \tag{17}$$

where p_{jt} is the average price for the brand in a county and $\frac{\partial \pi_{jt}}{\partial p_{jt}}$ is the derivative of the per-period profits. This derivative equals $\frac{\partial S_{jt}}{\partial p_{jt}}(p_{jt} - c_{jt}) + S_{jt}$ for single-brand firms. We measure marginal cost, c_{jt}, as the brand-specific wholesale rack price plus federal, state, and local gasoline taxes. Thus, unlike papers that infer marginal costs from firms' first-order conditions, our supply-side exercises use a directly observed measure of marginal costs.

The estimated dynamic parameters, along with marginal costs, allow us to compute the static profit derivative $\frac{\partial \pi_{jt}}{\partial p_{jt}}$ directly from the data. If this were zero, it would imply that firms are pricing myopically, as they are simply maximizing current-period profits. When it is nonzero, it implies that dynamic considerations, complementary profits, or other factors are affecting a firm's pricing decision.

On average, we find that $\frac{\partial \pi_{jt}}{\partial p_{jt}}$ is positive, implying that firms typically price below the static optimum. This is consistent with forward-looking behavior and the presence of dynamics, complementing the reduced-form findings of Section 3.3. However, other mechanisms can also generate a positive static-profit derivative. Gas stations may earn profits from convenience-store sales and other in-store services, which can rationalize below-static-optimum or even below-cost fuel pricing as part of a loss-leader strategy (see, for example, Houde, 2012; Chaves et al.,

Table 5: Summary of Implied $\beta \frac{\partial E[V_j(\cdot)|\cdot]}{\partial p_{jt}}$

Group	Mean	Min	p25	Median	p75	Max
All	-0.122	-0.715	-0.165	-0.094	-0.051	0.040

Notes: Table displays the estimated derivative of continuation value. A finding of zero would indicate the absence of forward-looking behavior by firms. Negative values indicate that firms are pricing lower in that period than the optimal myopic price.

2025). For the maintained model, equation (17) implies that the gap between the static first-order condition and zero is the derivative of the continuation value (DCV), $\beta \frac{\partial E[V_j(\cdot)|\cdot]}{\partial p_{jt}}$. Thus, the dynamic incentive is the residual that rationalizes the observed pricing behavior of the firms, conditional on the demand-side assumptions, the data, and Bertrand price competition. After estimating demand in an independent step, we recover this wedge directly.

Summary statistics for the value of the derivative of the continuation value (DCV) are presented in Table 5. The mean and median are negative, which implies that, typically, a reduction in price would increase the expected future return. We estimate a positive residual in only 3 percent of observations. The magnitudes are meaningful: the mean of -0.122 implies that a 1-cent increase in price would increase static profits by roughly 4 percent.[20] Under our dynamic-pricing interpretation, firms are reducing prices to invest in future demand. Such behavior allows firms to occasionally have negative margins, which occur in 2.7 percent of the observations in our data. This result, combined with our reduced-form findings of anticipatory pricing for expected costs, provides consistent evidence of forward-looking pricing behavior in retail gasoline.

5.2 Modeling Dynamic Incentives

To estimate counterfactual pricing behavior by firms, it is necessary to estimate how dynamic incentives vary with state variables and firm actions. To do so, we develop an approach that relies on the structural demand model to calculate the static component of profits, and we use a functional approximation to capture dynamic incentives. This provides a flexible empirical model that can incorporate time-varying demand and cost shocks.

Specifically, we approximate the dynamic component of firms' first-order conditions (the DCV) directly with a reduced-form model that is a function of state variables. Using the data and the estimated demand parameters, we obtain estimates of the DCV and project these estimates on state variables, including measures that capture expectations. We estimate the fol-

[20] The average profit in our scaled units, measured as margin times market share, is 0.029. Margins are approximately 21 cents per gallon.

lowing dynamic first-order condition:

$$\frac{\partial \pi_{jt}}{\partial p_{jt}} + \Psi_j(p_t, r_t, c_t, x_t; \theta) + \zeta_{jt} = 0. \quad (18)$$

For any observed or counterfactual data point, we construct $\frac{\partial \pi_{jt}}{\partial p_{jt}}$ directly using the structural demand estimates. We use $\Psi_j(\cdot)$ to approximate $\beta \frac{\partial E[V_j(\cdot)|\cdot]}{\partial p_{jt}}$ from equation (17), and ζ_{jt} is the unobserved error. We can use this function to approximate how the dynamic incentives change with the state and the endogenous pricing decisions by firms, allowing for counterfactual analysis. In general, Markovian assumptions allow for the continuation value to be expressed as a function of the state and firm actions.

This approach is an alternative to that of Bajari et al. (2007), who use an approximation to the policy function and, based on this, leverage model structure to estimate the dynamic incentives and static parameters. Conversely, we use structural modeling to obtain static parameters and calculate a reduced-form approximation to the dynamic incentives. Our approach has three advantages. First, the static component of profits is obtained without having to make any assumptions about firm expectations and discount rates. Second, we avoid the need to make dimension-reducing assumptions, such as constructing a limited grid for prices, that are less palatable in our setting.[21] Third, utilizing an approximation for the DCV greatly reduces the computational time needed to recompute equilibria. For our approach to accurately represent behavior, we require that the state variables included in the reduced-form approximation capture the payoff-relevant states (including market structure) and also that the counterfactual states can be reasonably interpolated from the data.

This reduced-form approach is consistent with a structural model (and solving for the equilibrium DCV) under the assumption that (i) the information set of firms matches the information set of the econometrician and (ii) firms perform limited forecasts of the evolution of future profits, consistent with the approximation used in estimation. Thus, our approach can be considered as an approximation to the forecasting behavior of firms that use simple regression-based forecasts of future profits. When this is the case, the firms' beliefs can correspond to the econometrician's estimates. To provide a sense of how close our estimates come to rational expectations, we use forward simulations to calculate realized profits when firms set prices according to equation (18). We discuss these forward simulations below.

To estimate $\Psi_j(\cdot)$, we project the estimated DCV onto the stock of affiliated consumers ($r_{jmt} \times \lambda_m$), the derivative of own share with respect to price, marginal costs, and expectations of future costs. Our model and the descriptive evidence suggest that these variables play an important role in expectations of future profits. We also include the fraction of state-dependent consumers λ_m, the number of stations, the total number of stations for all brands, and the

[21]Under the alternative policy function approach, an insufficiently flexible policy function may be incompatible with equilibrium prices.

Table 6: Dynamic Pricing Incentive: Regressions

	$\beta\frac{\partial E[V_j(\cdot)\mid\cdot]}{\partial p_{jt}}$ (1)	Sensitivity (2)
$r_{jmt} \times \lambda_m$	-0.889^{***}	6.428^{***}
	(0.001)	(0.017)
$-\frac{dS_{jmt}}{dp_{jmt}}$	-0.540^{***}	6.416^{***}
	(0.001)	(0.037)
Marginal Cost	-0.001^{***}	0.019^{***}
	(0.000)	(0.002)
Cost Change (30-Day Ahead)	0.014^{***}	-0.092^{***}
	(0.000)	(0.012)
λ_m	0.010^{***}	0.438^{***}
	(0.001)	(0.015)
Num. Stations (Brand)	0.000^{***}	0.011^{***}
	(0.000)	(0.000)
Num. Stations (Market)	-0.000^{***}	-0.001^{***}
	(0.000)	(0.000)
Num. Brands (Market)	-0.001^{***}	0.042^{***}
	(0.000)	(0.001)
Constant	Yes	Yes
Observations	110,844	110,844
R^2	0.973	0.785

Notes: Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Table displays the estimated coefficients from a regression of the dynamic pricing incentive on state variables. The second column reports the regression with a measure of sensitivity, which is the log absolute value of the dynamic pricing incentive. In general, a negative coefficient in the first column implies a greater sensitivity to dynamics when pricing, generating a positive coefficient in the second column.

number of brands as market-level controls. We do not include market-level or brand-level fixed effects. Instead, we use cross-market variation to quantify the relationships between the selected covariates and the DCV.

The results of estimating equation (18) are reported in Table 6. The first specification reports the coefficients from a regression of the DCV on the covariates. As the DCV is negative on average, a negative coefficient implies that the variable is associated with a stronger dynamic pricing incentive, or a greater deviation from the optimal static price. We flip the sign on the own-price derivative, which is also negative, to facilitate interpretation. The eight-parameter model has an R^2 of 0.97. Overall, the reduced-form approach captures the vast majority of the price variation that cannot be explained by static optimization. Because the dependent variable is itself recovered from the observed pricing condition using the same demand system, this fit should not be interpreted as an out-of-sample validation exercise. The high explanatory power of this parsimonious model nonetheless provides some confidence for interpolation and extrapolation in our counterfactual analysis.

We find that a higher share of affiliated consumers, $r_{jmt} \times \lambda_m$, increases the magnitude of the DCV, corresponding to an increased investment incentive when pricing. Thus, the presence of less-elastic affiliated consumers provides a dynamic incentive to keep prices low, even though they provide a direct incentive to raise static prices. The coefficient on the own-price derivative indicates that the dynamic incentive is greater when the derivative (which is also negative) is larger in magnitude. This indicates that demand from unaffiliated consumers also plays a role in dynamic incentives. One mechanism to explain this is that some unaffiliated (state-dependent) consumers become affiliated with that brand in the future, so the derivative captures the possibility of attracting more future affiliated consumers. Increases in marginal costs tend to make firms more sensitive to dynamic profit considerations, while increases in future expected costs lead firms to place less weight on future profits. Finally, we find that the number of stations and the number of brands have relatively small coefficients, after controlling for the above factors.

To help interpret how sensitive firms are to dynamic considerations, the second column of Table 6 reports a regression where we replace the value of the DCV with the logged absolute value. These coefficients reflect the semi-elasticity for the magnitude of the dynamic incentive. A positive coefficient in the second column indicates that an increase in the variable makes a firm more sensitive to dynamic considerations, whereas a negative coefficient indicates a reduced sensitivity to dynamic considerations when pricing. The results in the second column suggest that firms' dynamic considerations are most sensitive to the stock of affiliated consumers and the own-price derivative. The results show a modest marginal relationship between the magnitude of the dynamic incentive and the overall share of state-dependent consumers in the market.

Forward simulations reported in Appendix D.2 provide directional validation for this approximation: realized future-profit changes are positively correlated with the estimated DCV and are of similar magnitude.

5.3 Merger Simulations

We now apply the empirical model to evaluate the impact of dynamic pricing incentives on horizontal mergers. We simulate a merger between Marathon and BP, the largest and fourth-largest non-fringe brands by overall share in our sample. Out of 241 markets, they overlap in 75. In these markets, the average post-merger HHI is 1511, and the mean change in HHI is 383. In eight markets, the resulting HHIs exceed 2500 with changes greater than 200, meeting

Table 7: Merger Effects

Brand	Dynamic Model: Joint Pricing			Dynamic Model: Brand Consolidation			Static Model: Brand Consolidation		
	Price	Share	Profit	Price	Share	Profit	Price	Share	Profit
Marathon-BP	4.91	-18.09	25.26	2.37	10.92	51.20	5.35	-16.74	34.44
Other	-0.33	9.44	5.85	1.53	-9.60	2.95	0.61	3.86	11.28
Overall	1.55	-2.93	15.16	2.05	-0.38	26.08	2.40	-5.40	22.39

Notes: Table displays the mean percent changes in prices, shares, and profits from counterfactual mergers between two brands in our data. The first six columns provide estimates from dynamic models that account for consumer inertia. The first three columns report a counterfactual joint pricing merger and the next three columns report a counterfactual brand consolidation merger. The last three columns provide the estimates for a brand consolidation merger from a static model with the same linear utility specification estimated on the same data. The reported effects are averaged over the full 51 post-merger weeks, which include a transitional period to adjust to the new market structure. Price effects are weighted by shares.

typical thresholds presumed likely to enhance market power. The merger would change 12 markets from 3 firms to 2 firms and 18 markets from 4 firms to 3 firms. We allow the firms to merge at the beginning of September 2014 (the midpoint of our sample) and calculate counterfactual prices and shares for the following 51 weeks (periods 54–104 of the sample).[22]

We consider both joint pricing and brand consolidation scenarios. For joint pricing, the merged firm has pricing control over both brands, maintained as distinct entities. For brand consolidation, the merged firm consolidates assets under a single brand. For the joint pricing scenario, we must specify the cross-price effects on the continuation value, $\beta \partial E\left[V_k\right] / \partial p_{jt}$, when k and j are owned by the merged firm. For baseline results, we assume effects are proportional to diversion ratios D_{kj}, so $\partial E\left[V_k\right] / \partial p_{jt} = D_{jk} \partial E\left[V_k\right] / \partial p_{kt}$. The diversion ratios capture relative effects on shares, which are likely to correlate with effects on profits. Results are qualitatively similar with moderate changes in this scaling factor.

Table 7 displays the effects on prices, shares, and profits, averaging across the full 51 weeks post-merger.[23] The first three columns report effects from the joint pricing scenario, and the next three columns report effects from brand consolidation. The joint pricing scenario predicts that merging firms will raise prices by 4.9 percent, with an 18 percent decrease in shares and a meaningful increase in profits. Overall market prices increase by 1.6 percent, reflecting price changes by both merging and non-merging firms.

The brand consolidation scenario predicts a price increase for merging firms of 2.4 percent, roughly half of the joint pricing prediction. However, in brand consolidation, merging firms

[22]Because our inelastic affiliated customers might technically purchase at very high prices, we impose a choke price of $5 in demand and impose a penalty for prices exceeding this value. Across all merger counterfactuals, only eight observations approach the choke price.

[23]It takes approximately 10–15 weeks to converge to the longer-run price levels for the brand consolidation merger, while prices adjust almost immediately for the joint pricing merger. The reported effects include the transitional periods. We obtain similar price effects when excluding the first 25 post-merger weeks.

also realize an *increase* in market share. This occurs because brand consolidation provides the merging firm with an advantage in retaining affiliated consumers and yields a greater investment incentive. At pre-merger prices held fixed, the merging firm would accumulate greater shares over time due to superior ability to retain affiliated consumers. In the brand consolidation merger, the combined share of affiliated consumers (r) for merging firms rises from 0.39 in the observed baseline to 0.44; by contrast, in joint pricing, the average combined share falls to 0.32. This highlights how incentives to invest or harvest can vary across merger types.

Thus, a change in relative ability to retain affiliated consumers shifts incentives to invest or harvest, as shown by differences in behavior for merging firms across scenarios. Empirically, this is captured by the shift in static and dynamic components of the first-order condition. Joint pricing provides an immediate incentive for the merged firm to raise prices and increase current-period joint profits. Brand consolidation also has this static incentive, but there is also an immediate dynamic incentive to invest as the acquiring brand can now capture more state-dependent consumers, captured by the coefficient on $-\frac{dS_{jmt}}{dp_{jmt}}$ in Table 6. The relative strength of these forces in longer-run equilibrium determines which merger type leads to greater price changes.

Another key difference between the types of mergers is the effect on rival firms. In joint pricing mergers, rival firms see a small decrease in prices (on average), and a 9.4 percent gain in market share. By contrast, brand consolidation leads rivals to increase prices by 1.5 percent and lose share as consumers are attracted to the consolidated brand. The small negative rival price effect in the joint pricing counterfactual reflects the dynamic pricing incentives faced by rivals as the merger changes the distribution of affiliated consumers. More specifically, joint pricing induces the merged firm to raise prices and lose affiliated share. In turn, there are more potentially affiliated customers for non-merging firms to attract, which strengthens their incentive to price aggressively. This investment incentive slightly offsets the usual static incentive for rivals to raise prices when the merging firm increases prices. In the brand consolidation setting, the merging firms implement smaller price increases and gain share; as a result, there is a smaller pool of affiliated customers for non-merging firms, and the typical incentive to increase price in response to the merging firms dominates.

For comparison, we report results from a merger analysis using a static model in the last three columns of Table 7. We use the standard logit demand system estimated in column (3) of Table 3, which employs the same linear utility specification as the dynamic model, i.e., equation (15). We simulate a brand consolidation merger for illustration. The static model predicts price effects of 5.4 percent for merging firms—more than double the dynamic model's brand consolidation prediction. The median price change for merging firms is 4.6 percent in the static model. By comparison, roughly one-third of markets in the joint pricing counterfactual exceed that value, while only 9 out of 75 markets have price changes at least that large in the dynamic brand consolidation counterfactual.

We find that predictions of a dynamic model with consumer inertia can diverge substantially from those of a static model. The dynamic incentive to invest in future demand can mitigate the short-run incentive to raise prices post-merger, dampening the exercise of horizontal market power. Further, how the merger is implemented can have a large impact on equilibrium prices. In this case, brand consolidation provides merging firms with an added incentive to invest, leading to average price increases less than half those predicted by a static model.

At the same time, the ranking of markets is similar across the three counterfactuals. Across markets, the correlation between any two of the three merger scenarios ranges from 0.713 to 0.750, and greater combined pre-merger share is associated with larger post-merger price increases in all cases. The key difference is therefore not whether mergers between larger firms could lead to higher price changes, but rather how much prices move and whether consolidating the brands post-merger attenuates or amplifies the static effects.

For context, Appendix D.3 considers the role of dynamic market power—the ability to raise prices due to consumer inertia—for this empirical setting. We provide counterfactuals that isolate dynamic pricing effects by varying the prevalence and the strength of inertia while holding market structure fixed.

5.4 Exploring Dynamic Incentives and Merger Prediction Bias

We further explore the properties of our dynamic demand and supply model. The aim of the analysis is to provide additional insight into the interaction of consumer inertia and horizontal mergers, and its implications for antitrust enforcement. We focus on the steady-state equilibrium of the model, which is the setting most commonly analyzed when performing merger simulations as part of a merger investigation.[24] This approach abstracts from period-to-period variation in demand and cost shocks to focus on long-run effects.

We start by characterizing the steady-state equilibrium properties of the model in a deterministic setting where marginal costs are constant. This provides an analytic expression for merger incentives in a dynamic setting. We then use numerical methods to analyze steady-state prices and the effect of horizontal mergers. This analysis demonstrates how the model's key mechanisms—the interplay between investment and harvest incentives—operate in equilibrium. Furthermore, this analysis highlights that whether or not a merger results in brand consolidation has important ramifications for post-merger prices.

5.4.1 Steady-State Merger Effects

We again assume that firms face the demand setting developed in Section 2: consumer utility is given by equation (4), choice probabilities are determined by equation (2), and aggregate share is given by equation (5).

[24] See, for example, Miller and Sheu (2021).

Consider a deterministic steady state where marginal costs are constant. Each firm k sells a set of products $j \in J_k$ and maximizes the expected discounted value of profits, so firm k's value function takes the form:

$$V_k(r) = \max_{p_k|p_{-k}} \pi_k(p,r) + \beta V_k(r'), \tag{19}$$

where p and r are vectors of prices and affiliated customers, respectively, and $r' \equiv f(p,r)$ specifies the shares of affiliated customers for each product in the next period.[25] In accordance with the demand model, an element of r' is $r'_j = \frac{1}{\lambda}(S_{jt} - (1-\lambda)s_{jt}(0))$. Static profits are $\pi_k(p,r) = \sum_{j \in J_k}(p_j - c_j) \cdot s_j(p,r)$.

To find the steady-state prices and affiliated shares for each firm, we focus on Markov perfect equilibrium. Firm k's profit-maximizing first-order conditions with respect to prices are:

$$\frac{\partial \pi_k}{\partial p_j} + \beta \frac{dV_k(r')}{dr'} \frac{dr'}{dp_j} = 0 \quad \forall j \in J_k. \tag{20}$$

Next, we take derivatives of equation (19) with respect to r and evaluate them at the prices that solve each firm's first-order conditions, which are the prevailing steady-state prices. Then, in conjunction with the steady-state condition $\frac{dV'}{dr'} = \frac{dV}{dr}$, we obtain the following system of equations:

$$\underbrace{\frac{dV_k(r)}{dr}}_{J\times 1} = \underbrace{\left[\frac{\partial \pi_k}{\partial p}\frac{dp}{dr} + \frac{\partial \pi_k}{\partial r}\right]}_{J\times 1} \underbrace{\left[I - \beta f_p(p,r)\frac{dp}{dr} - \beta f_r(p,r)\right]^{-1}}_{J\times J}. \tag{21}$$

In this equation, $\frac{\partial \pi_k}{\partial p}$, $\frac{\partial \pi_k}{\partial r}$, $f_p(p,r)$, and $f_r(p,r)$ are known conditional on values of p and r. The remaining unknowns are $\frac{dp}{dr}$ and $\frac{dV_k(r)}{dr}$.

Consistent with Dubé et al. (2009), equilibrium prices may be increasing or decreasing in the level of affiliation. With greater affiliation, firms face less elastic demand, but the incentive to invest in future demand tends to increase. Whether prices increase with affiliation depends on the relative weights on these forces and underlying market parameters. Appendix E.5 establishes this ambiguity analytically in a monopoly benchmark. Importantly for our context, these forces interact with market structure, so that the presence of affiliation can have differential effects on prices post-merger.

Consider a merger in which firm k acquires product b and maintains it as a separate brand. In the case of static demand, the post-merger change in pricing incentives for a product $j \in J_k$ at pre-merger equilibrium prices is $\sum_{l \in \{J_k, b\}} \frac{\partial \pi_l}{\partial p_j} - \sum_{l \in J_k} \frac{\partial \pi_l}{\partial p_j} = \frac{\partial \pi_b}{\partial p_j}$, which is positive for substitutes, leading the firm to raise prices.

[25]Note that the value function has no expectations operator, as we consider a deterministic steady state in this subsection.

For dynamic demand, the post-merger change at the pre-merger steady-state prices is:

$$\textit{Change in dynamic first-order condition:} \quad \frac{\partial \pi_b}{\partial p_j} + \beta \left(\frac{d\tilde{V}_k(r')}{dr'} - \frac{dV_k(r')}{dr'} \right) \frac{dr'}{dp_j}, \tag{22}$$

where $\tilde{V}_k$ incorporates the discounted flow of profits from product b. The first term is equivalent to the change in the static first-order condition, though evaluated at different equilibrium prices. The second term may be positive or negative, so consumer dynamics can either exacerbate or mitigate incentives to raise prices post-merger. On the one hand, the acquiring firm internalizes that a price increase for product j will increase customers affiliated with the new brand b, increasing future profits. On the other hand, such a price increase would reduce the number of customers affiliated with brand j, indirectly incentivizing lower future prices for all firms and reducing future profits for the acquired brand. The trade-off can be conceptualized as balancing two forces: the recapture of lost affiliated customers by the acquired brand versus maintaining a higher stock of affiliated consumers to soften competition.

5.4.2 Numerical Analysis

We now use numerical simulations to illustrate more systematically how merger effects vary due to demand dynamics and the specifics of merger implementation. We analyze symmetric three-firm markets, with consumer indirect utility following equation (4).

Market parameters are drawn such that $\xi \in [0, 10]$, $\overline{\xi} \in [0, 10]$, and $\alpha \in [-10, 0]$, with marginal cost normalized to one. For each parameter draw, we construct markets for $\lambda \in \{0.05, 0.1, 0.15, \ldots, 0.70\}$. We impose the steady-state condition governing the evolution of affiliated customers, $r' = r$, and the system of equations in (21), allowing us to solve for steady-state prices and shares conditional on the derivative matrix $\frac{dp}{dr}$. We determine this matrix numerically using a local approximation method. After solving for steady-state outcomes, we restrict analysis to markets with individual firm shares between 0.05 and 0.30 and markups, defined as $\frac{p-c}{p}$, between 0.05 and 0.75, yielding 6,566 markets. Full details of the simulation methodology and descriptive statistics of our simulations are provided in Appendix E. Appendix E.3 summarizes how the pricing effects of consumer inertia and mergers vary across simulated markets, and Appendix E.4 reports representative example markets.

To show the implications for mergers, we analyze a joint pricing and a brand consolidation merger that combines two of the single-product firms in each simulated market. We also consider the following scenario: a practitioner observes each firm's pre-merger prices, marginal costs, and aggregate market shares. These data are then used to recover the demand parameters of the standard logit model (without affiliation), and then the price effects of a merger are simulated. Because we have symmetric firms, the static logit merger predictions are identical for joint pricing and brand consolidation mergers (see Section 2.2). We compare the results of

Figure 4: Price Increase by Market Share

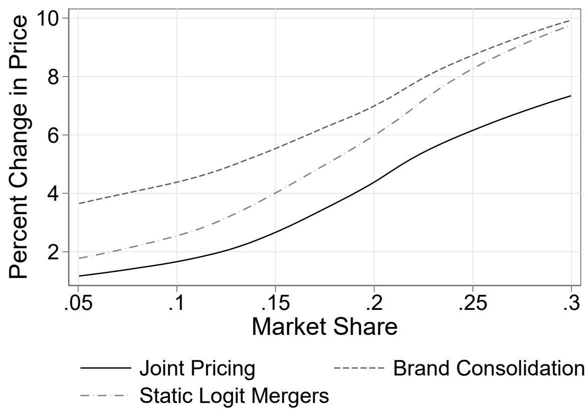

Notes: The lines depict local polynomial regressions of the merged firm's percentage price change on its pre-merger market share.

this misspecified simulation to the results using the correct underlying dynamic demand model.

The "average" market in our simulations is one that would typically raise moderate concern from the U.S. antitrust agencies, should a merger occur. The mean HHI (1067) falls in the "unconcentrated" range, but the change in HHI (712) generally warrants a thorough investigation. The average pre-merger difference between price and cost is 0.26, and the mean market share is 0.17. These mean values imply an "Upward Pricing Pressure" index of $\frac{0.17}{1-0.17} \cdot (1.26-1) = 0.053$, which is just over the threshold that may trigger an investigation.[26] The full range of markets spans those that would receive no scrutiny and those that almost certainly would be challenged. Thus, the simulations generate a reasonable set of markets within which to explore merger effects.

Figure 4 plots the merger price effects related to the pre-merger market share of each symmetric firm. To generate the graph, we run a local polynomial regression of the merger price effect on one of the symmetric firm's pre-merger market share. We generate fitted lines for (i) the joint pricing merger effect, (ii) the brand consolidation merger price effect, and (iii) the misspecified static model. In line with intuition, the price effect of a merger is increasing in pre-merger market shares. The average price increase is 3.8 percent for joint pricing mergers and 6.5 percent for brand consolidation mergers. Note that, in contrast to our empirical application, the brand consolidation merger leads to larger price effects on average. The ordering can flip with changes in primitives and market structure.

In these simulations, the average price increase from the static model falls between the two, with an average of 5 percent. For joint pricing mergers, the static prediction is less biased when pre-merger shares are large. By contrast, for brand consolidation mergers, the static prediction

[26]See, for example, Farrell and Shapiro (2010) and Miller et al. (2017). This calculation assumes that diversion is proportional to market share, which is often assumed at the early stages of an antitrust investigation.

Figure 5: Simulations: Merger Prediction Bias

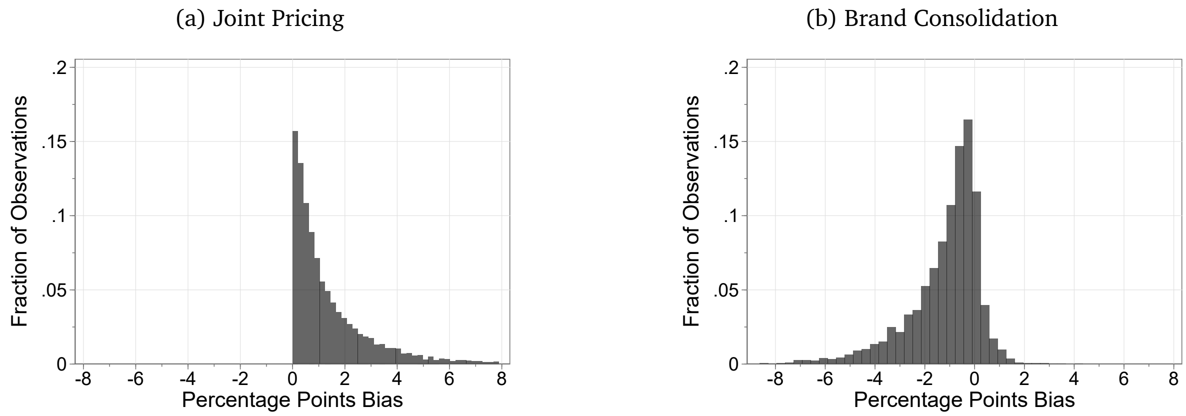

Notes: Prediction bias is defined as the prediction of a (misspecified) static logit model minus the true price increase. Panel (a) depicts the bias when the merger consolidates pricing control of two products. Panel (b) depicts the bias when the merger consolidates two products under one brand. A brand consolidation merger is defined in the main text.

is less biased when pre-merger shares are small.

Figure 5 plots the bias, defined as the (incorrect) static prediction minus the (correct) dynamic prediction, across all 6,566 markets. Panel (a) depicts the distribution of bias for joint pricing mergers. For the included markets, the static model overpredicts the true dynamic effect, with a mean bias of 67.3 percent (when scaled by the magnitude of the true effect). This overprediction is not a guarantee of the model; in markets that did not meet our inclusion criteria, the static model underpredicted the dynamic effect. Further, asymmetries across firms can also yield bias in the other direction. The figure shows that the bias in brand consolidation mergers may be positive or negative. On average, we find a bias of -19.7 percent. While static models are more likely to underpredict the true dynamic price effect, almost 25 percent of simulations result in overpredictions.

In both types of mergers, biases arise primarily from the omission of dynamic incentives to invest in future demand rather than from biased elasticity estimates alone.[27] Figure 5 shows how, holding fixed the same misspecified elasticity, we can obtain biases of the opposite sign (positive or negative), depending on the merger type. This shows that consumer dynamics matter for merger simulations, as markups and elasticities alone are not sufficient when firms care about future demand.

[27]The static model yields more elastic demand, on average, than the dynamic model in these simulations (elasticity of -4.74 vs. -3.86). More elastic demand generates smaller merger price effects in the logit model.

5.5 Discussion and Implications

Our analysis demonstrates that failing to account for consumer dynamics in merger simulation can lead to materially biased predictions. The size, and even the direction, of the bias depends on how the merger is implemented and on how the new firm restructures brands. Static models miss changes in dynamic investment and harvest incentives, so two mergers with similar pre-merger observables can generate very different post-merger price effects once consumer inertia is taken into account.

These findings have practical implications for antitrust enforcement. Antitrust agencies often infer elasticities from markups calculated using accounting data, which omit firms' dynamic incentives (see Miller et al., 2013). In addition to generating potentially incorrect elasticities, failing to account for dynamic incentives in first-order conditions can have large direct effects on post-merger predictions. Our results suggest that a careful assessment of how a proposed merger will be implemented—whether the merging parties will maintain separate brands or consolidate—may be particularly important in markets characterized by consumer inertia.

More broadly, our framework highlights the importance of incorporating consumer dynamics when evaluating market power. The empirical model provides a tractable approach to quantifying the interplay between investment and harvest incentives that determine equilibrium pricing. By combining structural demand estimation with a reduced-form approximation to dynamic incentives, we can assess both horizontal market power (from changes in market structure) and dynamic market power (from consumer inertia) within a unified framework. The methodology can be applied to other settings where switching costs, loyalty, or habit formation create persistent consumer-firm relationships that affect competitive dynamics.

6 Conclusion

We develop a tractable model of consumer inertia for merger analysis. The framework captures dynamic pricing incentives generated by habit, loyalty, or switching frictions while remaining estimable with market-level data. We show that accounting for consumer inertia is important when performing counterfactual exercises, such as merger simulation.

Using retail gasoline data, we document reduced-form evidence consistent with inertia on both the demand and supply sides. Consumers' purchase histories display strong repeat buying and weaker persistence after longer purchase gaps. On the supply side, new retail locations initially price below incumbents and then raise prices over time, while stations also adjust prices gradually in response to predictable cost changes. Even in this relatively competitive setting with a fairly homogeneous product, dynamic behavior is economically important.

We develop and estimate an empirical model that can identify dynamic demand parameters using data on price, shares, and an instrument. The structural estimates imply that 63 percent of consumers are prone to affiliation and that affiliated consumers are relatively insensitive to price. Unaffiliated consumers are price sensitive and play an important role in disciplining equilibrium prices. We show, both theoretically and empirically, that static merger simulations can be materially biased because they omit dynamic terms in firms' first-order conditions. They can also miss an important institutional margin: whether the merged firm keeps separate brands under joint pricing or consolidates them into a single brand.

References

BAJARI, P., C. L. BENKARD, AND J. LEVIN (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75, 1331–1370.

BEGGS, A. AND P. KLEMPERER (1992): "Multi-Period Competition with Switching Costs," *Econometrica*, 60, 651.

BERGEMANN, D. AND J. VÄLIMÄKI (1996): "Learning and Strategic Pricing," *Econometrica*, 64, 1125.

——— (2006): "Dynamic Price Competition," *Journal of Economic Theory*, 127, 232–263.

BERRY, S., M. CARNALL, AND P. SPILLER (2006): "Airline Hubs: Costs, Markups and the Implications of Customer Heterogeneity," *Advances in Airline Economics*, 1.

BERRY, S. AND P. JIA (2010): "Tracing the Woes: An Empirical Analysis of the Airline Industry," *American Economic Journal: Microeconomics*, 2, 1–43.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841.

BESANKO, D., U. DORASZELSKI, AND Y. KRYUKOV (2018): "How Efficient Is Dynamic Competition? The Case of Price as Investment," NBER Working Paper 23829, Cambridge, MA.

BRONNENBERG, B. J., J.-P. DUBÉ, AND S. MOORTHY (2019): "The economics of brands and branding," in *Handbook of the Economics of Marketing*, Elsevier, vol. 1, 291–358.

BRONNENBERG, B. J., J.-P. H. DUBÉ, AND M. GENTZKOW (2012): "The Evolution of Brand Preferences: Evidence from Consumer Migration," *American Economic Review*, 102, 2472–2508.

CABRAL, L. (2011): "Dynamic Price Competition with Network Effects," *The Review of Economic Studies*, 78, 83–111.

CARRANZA, J. E. (2010): "Product Innovation and Adoption in Market Equilibrium: The Case of Digital Cameras," *International Journal of Industrial Organization*, 28, 604–618.

CHANDRA, A. AND M. TAPPATA (2011): "Consumer Search and Dynamic Price Dispersion: An Application to Gasoline Markets," *The RAND Journal of Economics*, 42, 681–704.

CHAVES, D., R. CLARK, M. DUARTE, AND J.-F. HOUDE (2025): "Misallocation in Differentiated Product Markets: Evidence from Collusion under Uniform Pricing," Working paper.

DONNA, J. D. (2021): "Measuring Long-Run Gasoline Price Elasticities In Urban Travel Demand," *RAND Journal of Economics*, 52, 945–994.

DUBÉ, J.-P., G. J. HITSCH, AND P. E. ROSSI (2009): "Do Switching Costs Make Markets Less Competitive?" *Journal of Marketing Research*, 46, 435–445.

——— (2010): "State Dependence and Alternative Explanations for Consumer Inertia," *The RAND Journal of Economics*, 41, 417–445.

EIZENBERG, A. AND A. SALVO (2015): "The Rise of Fringe Competitors in the Wake of an Emerging Middle Class: An Empirical Analysis," *American Economic Journal: Applied Economics*, 7, 85–122.

FARRELL, J. AND P. KLEMPERER (2007): "Chapter 31 Coordination and Lock-In: Competition with Switching Costs and Network Effects," in *Handbook of Industrial Organization*, Elsevier, vol. 3, 1967–2072.

FARRELL, J. AND C. SHAPIRO (1988): "Dynamic Competition with Switching Costs," *The RAND Journal of Economics*, 19, 123.

——— (2010): "Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition," *The BE Journal of Theoretical Economics*, 10.

GANDHI, A. AND J.-F. HOUDE (2019): "Measuring Substitution Patterns in Differentiated-Products Industries," NBER Working Paper 26375.

GOWRISANKARAN, G. AND M. RYSMAN (2012): "Dynamics of Consumer Demand for New Durable Goods," *Journal of Political Economy*, 120, 1173–1219.

HANDEL, B. R. (2013): "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *The American Economic Review*, 103, 2643–2682.

HENDEL, I. AND A. NEVO (2013): "Intertemporal Price Discrimination in Storable Goods Markets," *American Economic Review*, 103, 2722–2751.

HO, C.-Y. (2015): "Switching cost and deposit demand in China," *International Economic Review*, 56, 723–749.

HO, C.-Y., M. RYSMAN, AND Y. WANG (2025): "Demand for Performance goods: Import Quotas in the Chinese movie market," Working paper.

HONKA, E. (2014): "Quantifying Search and Switching Costs in the US Auto Insurance Industry," *The RAND Journal of Economics*, 45, 847–884.

HORTAÇSU, A., S. A. MADANIZADEH, AND S. L. PULLER (2017): "Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market," *American Economic Journal: Economic Policy*, 9, 192–226.

HOUDE, J.-F. (2012): "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline," *American Economic Review*, 102, 2147–2182.

KEANE, M. P. (1997): "Modeling Heterogeneity and State Dependence in Consumer Choice Behavior," *Journal of Business & Economic Statistics*, 15, 310–327.

LEE, R. S. (2013): "Vertical Integration and Exclusivity in Platform and Two-Sided Markets," *American Economic Review*, 103, 2960–3000.

LEVIN, L., M. S. LEWIS, AND F. A. WOLAK (2017): "High Frequency Evidence on the Demand for Gasoline," *American Economic Journal: Economic Policy*, 9, 314–347.

LEWIS, M. AND M. NOEL (2011): "The Speed of Gasoline Price Response in Markets with and without Edgeworth Cycles," *Review of Economics and Statistics*, 93, 672–682.

MELNIKOV, O. (2000): "Demand for Differentiated Durable Products: The Case of the U.S. Computer Printer Market," Working paper, Department of Economics, Yale University.

MILLER, N., M. REMER, C. RYAN, AND G. SHEU (2017): "Upward Pricing Pressure as a Predictor of Merger Price Effects," *International Journal of Industrial Organization*, 52, 216–247.

MILLER, N. H., M. REMER, AND G. SHEU (2013): "Using Cost Pass-through to Calibrate Demand," *Economics Letters*, 118, 451–454.

MILLER, N. H. AND G. SHEU (2021): "Quantitative Methods for Evaluating the Unilateral Effects of Mergers," *Review of Industrial Organization*, 58, 143–177.

NISHIDA, M. AND M. REMER (2015): "The Determinants and Consequences of Search Cost Heterogeneity: Evidence from Local Gasoline Markets," SSRN Working Paper 2628576.

PAKES, A., M. OSTROVSKY, AND S. BERRY (2007): "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)," *The RAND Journal of Economics*, 38, 373–399.

REMER, M. (2015): "An Empirical Investigation of the Determinants of Asymmetric Pricing," *International Journal of Industrial Organization*, 42, 46–56.

REMER, M. AND F. R. WARREN-BOULTON (2014): "United States v. H&R Block: Market Definition in Court since the 2010 Merger Guidelines," *The Antitrust Bulletin*, 59, 599–618.

SHCHERBAKOV, O. (2016): "Measuring Consumer Switching Costs in the Television Industry," *The RAND Journal of Economics*, 47, 366–393.

SHUM, M. (2004): "Does Advertising Overcome Brand Loyalty? Evidence from the Breakfast-Cereals Market," *Journal of Economics & Management Strategy*, 13, 241–272.

STAHL, D. O. (1989): "Oligopolistic Pricing with Sequential Consumer Search," *The American Economic Review*, 700–712.

VALLETTI, T. AND H. ZENGER (2021): "Mergers with Differentiated Products: Where Do We Stand?" *Review of Industrial Organization*, 59, 179–212.

Appendix

A Brand Consolidation and Joint Pricing Mergers

A.1 Implementation of Joint Pricing Mergers

The implementation of joint pricing mergers is straightforward. We maintain the utility parameters from the demand side, and we assign the merged brands to the same firm. The merged entity sets prices for both brands, internalizing the cross-price effects on the joint profits.

A.2 Implementation of Brand Consolidation Mergers

For a brand consolidation merger, we remove one brand (the acquired brand) from the market. For this type of merger, there is a question of how the acquired assets translate into demand for the merging firm.

In this paper, we assume that the shares of unaffiliated consumers (shoppers) would remain the same if prices were maintained at the same level. Implicitly, we assume that the merged entity retains the same retail locations, and that the brand has no effect on shoppers.

To implement this, we adjust the demand shock for the remaining product of the merged firm so that the choice probabilities of unaffiliated customers are unchanged at pre-merger prices. Let a denote the acquiring brand, and b denote the acquired (and removed) brand. Recall that the choice probabilities for unaffiliated consumers are given by

$$s_{jt}(0) = \frac{\exp(\delta_{jt})}{1 + \sum_k \exp(\delta_{kt})} \tag{A.1}$$

yielding $\delta_{jt} = \ln(s_{jt}(0)/s_{0t}(0))$. Recall that $\delta_{jt} = \xi_{jt} + \alpha p_{jt}$. For the merged brand, we adjust the utility-shock fixed effect to ξ'_{at}, where

$$\xi'_{at} = \ln\left(\frac{s_{at}(0) + s_{bt}(0)}{s_{0t}(0)}\right) - \alpha \overline{p}_{ab} \tag{A.2}$$

Here, $s_{at}(0) + s_{bt}(0)$ is the combined pre-merger unaffiliated share of products a and b, and $\overline{p}_{ab}$ is the share-weighted average price of products a and b. For non-merging firms, we maintain $\xi'_{jt} = \xi_{jt}$.

This adjustment ensures that, if prices are held fixed (including share-weighted prices for the merging brand), the choice probabilities of shoppers for all products, including the non-merging products, are unchanged. Then, given the adjustment, we allow the firms to price optimally.

Under these assumptions, the merging firm gets some benefit for affiliated consumers. Our adjustment implies that $\exp(\delta'_{at}) = \exp(\delta_{at}) + \exp(\delta_{bt})$, and therefore $\exp(\delta'_{at} + \overline{\xi}) = \exp(\delta_{at} + \overline{\xi}) + \exp(\delta_{bt} + \overline{\xi})$. Making the appropriate adjustment to the choice probability equations, this implies that consumers affiliated with the merged brand choose it with the same probability as if they were affiliated with both brands a and b pre-merger. In implementation, we assume that consumers who were affiliated with brand b become unaffiliated after the merger. We found that whether or not affiliation transferred to the acquiring brand made little difference for our counterfactuals.

A.3 Equivalence of Joint Pricing and Brand Consolidation in Symmetric Logit

We now prove that in a static logit model, a joint pricing and brand consolidation merger (as defined in the previous subsection) will produce the same price effect if the merging firms are symmetric.

Lemma 1: Suppose the following is true,

(i) Demand is characterized by standard logit.

(ii) While holding all else equal, at a price $\bar{p}$, splitting firm m into two firms j and k yields equal shares that sum to the share of the original firm: $s_m(\bar{p}) = 2s_j(\bar{p}) = 2s_k(\bar{p})$.

It follows that $s_m(p) = 2s_j(p) \;\; \forall p$, i.e., the relation in (ii) holds for any price.

Proof: By construction:

$$s_m = \frac{e^{\xi_m + \alpha\bar{p}}}{1 + e^{\xi_m + \alpha\bar{p}} + \sum_g e^{\xi_g + \alpha p_g}} = 2\frac{e^{\xi_j + \alpha\bar{p}}}{1 + 2e^{\xi_j + \alpha\bar{p}} + \sum_g e^{\xi_g + \alpha p_g}} = 2s_j \qquad \text{(A.3)}$$

The equality above holds if and only if $e^{\xi_m + \alpha\bar{p}} = 2e^{\xi_j + \alpha\bar{p}}$. Dividing both sides by $e^{\xi_j + \alpha\bar{p}}$ and then taking logs, we have $\xi_m + \alpha\bar{p} - \xi_j - \alpha\bar{p} = \log(2)$. Therefore, $\xi_m - \xi_j = \log(2)$. It follows that if (i) and (ii) are true, then $s_m(p) = 2s_j(p) \;\; \forall p$.

Lemma 2: Suppose demand is characterized by logit. Suppose a single-product firm has marginal cost c, market share s_m, and profit-maximizing price p^*. Then a two-product firm with marginal cost c and product market shares given by $s_j(p) = \frac{s_m(p)}{2}$ will set the same profit-maximizing price p^*.

Proof: With logit demand, the first-order condition of one product for an N product firm is:

$$\frac{d\Pi}{dp_1} = 1 + \alpha(p_1 - c_1)(1 - s_1) - \alpha\sum_{l=2}^{N}(p_l - c_l)s_l = 0 \qquad \text{(A.4)}$$

Now, suppose all of the firm's products have the same marginal cost, c, and that all of its products are symmetric, $\delta_n = \delta$ for all n. Then equation (A.4) simplifies to:

$$\frac{d\Pi}{dp_n} = 1 - \alpha(p-c)(s-1) - \alpha(N-1)(p-c)s = 0 \tag{A.5}$$

Solving equation (A.5) for p yields the symmetric, profit-maximizing price for each product:

$$p^* = -\frac{1}{\alpha}\left[\frac{1}{1-Ns}\right] + c \tag{A.6}$$

Now, suppose a single-product firm sets a profit-maximizing price of p^* and therefore has market share s_m. Also hold the number and characteristics of all other firms in the market constant. Now suppose we replace the single-product firm with a two-product firm with marginal cost c and product characteristics δ_j such that $s_j(p^*) = \frac{s_m(p^*)}{2}$. By equation (A.6), the two-product firm will set the same profit-maximizing price, p^*. We therefore prove Lemma 2.

Lemma 1 and Lemma 2 help prove the following proposition.

Proposition: Let demand be characterized by logit. Consider the following two mergers of symmetric single-product firms with marginal cost c:

(i) *Joint pricing*: After the merger, the firm retains both products and prices them to jointly maximize post-merger profits.

(ii) *Brand consolidation*: After the merger, the firm removes one of its brands from the market. For the consolidated product, the post-merger share would have the same market share as the sum of the two single-product pre-merger firms at pre-merger prices. Post-merger, the firm sets a profit-maximizing price for the one product it keeps in the market.

The post-merger price for the two products in merger (i) is the same as the price for the one remaining product in merger (ii).

B Additional Reduced-Form Evidence on Dynamics

B.1 Dynamic Demand: Repeat Visits and Purchase Intervals

We build on our analysis from Section 3.2 using the NielsenIQ Consumer Panel Data. Here, we leverage the data on the duration between purchases to provide additional evidence about repeat purchases and consumer inertia. In our model, consumers who choose the outside option lose their affiliation to a particular brand. In the data, choosing the outside option is captured by longer periods between purchases, when the consumer chooses not to purchase. We denote the number of days since the previous purchase as the "purchase interval." Longer intervals could arise due to idiosyncratic, household-specific shocks, or, for example, higher prices.

We regress an indicator for repeat purchase on the purchase interval while controlling for household-year and weekly fixed effects. The estimates are reported in column (1) of Table B.1. We again find evidence consistent with our model of consumer inertia. In particular, the coefficient of -0.0011 indicates that a longer purchase interval is negatively related to the probability of a repeat purchase. If households delay for a week (i.e., choose the outside good), they are roughly 1% less likely to return to the same brand. This again suggests that past decisions affect brand choice and is consistent with our model in which consumers lose their affiliation when choosing the outside good.

To provide more detail on these dynamics, we subset households by the fraction of their purchases that are repeat purchases, based on the data displayed in Figure 1. Bin 1 households have less than 25% repeat purchases. Bin 2 includes those whose fraction is between 25% and 50%, Bin 3 includes those between 50% and 75%, and Bin 4 includes those between 75% and 99%. For this exercise, we exclude households with 100% repeat visits. These buckets are used to proxy for consumers who may be more or less prone to inertia.

We then repeat the regression exercise separately for each bin. The estimates are reported in columns (2)–(5) of Table B.1. The estimates indicate that the propensity of Bin 1 consumers to have a repeat purchase is not affected by the purchase interval. This is consistent with our assumption that some consumers are "shoppers" and are unaffected by inertia. The relationship between the purchase interval and the probability of a repeat purchase is stronger in columns (3), (4), and (5), which may reflect the fact that a greater proportion of households in those bins are prone to inertia/affiliation.

Finally, as a placebo test, we have also run a regression in which we replace the dependent variable (the interval from the previous purchase) with the interval until the next *future* purchase. The coefficient estimate is almost exactly 0, with a t-stat of -0.01. This provides some assurance that results in column (1) are due to consumer behavior and not spurious trends in the data.

Table B.1: Predicting a Repeat Visit

	(1)	(2)	(3)	(4)	(5)
Purchase Interval	-0.0011***	0.0001	-0.0006***	-0.0015***	-0.0012***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Constant	0.7329***	0.2060***	0.4054***	0.6425***	0.9285***
	(0.000)	(0.003)	(0.001)	(0.001)	(0.000)
Observations	2680548	42540	524061	766642	1347305
Household-Year FE	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes
Household-Year Group	All	Bin 1	Bin 2	Bin 3	Bin 4

Notes: The dependent variable is an indicator of whether a purchase is a repeat visit. Bin 1 households are those whose fraction of repeat purchases is less than 25% of their total purchases. Bin 2 includes those whose fraction is between 25% and 50%, Bin 3 includes those between 50% and 75%, and Bin 4 includes those between 75% and 99%. The sample is restricted to households with at least 26 gasoline purchases in a year and purchase intervals of less than 60 days.

B.2 Dynamic Demand: Correlation in Shares Over Time

Though ultimately the importance of demand-side dynamics in the data is estimated by the model, it is informative to examine the reduced-form relationships between key elements. The dynamic model is one in which today's quantity depends on the quantity sold last period. As motivation for this model, we present the results from reduced-form regressions of shares on lagged shares in Table B.2.

This exercise demonstrates that even after including rich fixed effects to capture static variation in consumer preferences, lagged shares are a significant predictor of current shares. The residual correlation in shares over time in the most detailed specification captures deviations from specific county-brand seasonal patterns. A positive correlation is consistent with state dependence in consumption. In specification (2), we show that lagged shares explain 95 percent of the variance in current shares, and the coefficient is close to one. In specification (3), we include measures of competition in the regressions, as well as a second-order polynomial in own price. The competition measures, which include the mean and standard deviations of competitor prices, are correlated with shares, but lagged shares still are the most important predictor of current shares. In specification (4), we include time and brand-county fixed effects. In the final specification (5), we include rich multi-level fixed effects: county-brand-(week of year), brand-state-week, and week-county. The coefficient of 0.628 on lagged shares in this specification indicates that deviations in shares are highly correlated over time, even when we condition on the most salient variables that would appear in a static analysis, adjust for brand-county specific seasonal patterns, and allow for flexible brand-state and county time trends. This finding is consistent with demand-side dynamics, as there are patterns in shares over time

Table B.2: Regressions with Share as the Dependent Variable

	(1)	(2)	(3)	(4)	(5)
Price	0.009***	0.000**	0.004	−0.005	−0.073***
	(0.001)	(0.000)	(0.002)	(0.004)	(0.017)
Lagged Share		0.973***	0.963***	0.554***	0.628***
		(0.001)	(0.001)	(0.002)	(0.003)
Price Squared			−0.000	0.001	0.010***
			(0.000)	(0.001)	(0.003)
Comp. Price (Mean)			−0.004***	−0.002	−0.108**
			(0.001)	(0.001)	(0.045)
Comp. Price (SD)			−0.000	0.003*	0.086***
			(0.001)	(0.001)	(0.024)
Comp. Stations			−0.000***	−0.000***	−0.004***
			(0.000)	(0.000)	(0.000)
Num. Stations			0.000***	0.004***	
			(0.000)	(0.000)	
Num. Brands			−0.001***	−0.003***	
			(0.000)	(0.000)	
Week FEs				X	
County-Brand FEs				X	
Brand-State-Week FEs					X
Week-County FEs					X
County-Brand-WofY FEs					X
Observations	174421	169931	169788	169770	156078
R^2	0.00	0.95	0.95	0.96	0.98

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

that are challenging to explain with contemporaneous variables.[28]

B.3 Cost Pass-through: Identifying Expected and Unexpected Costs

We now analyze gas stations' dynamic reactions to expected and unexpected costs. To disentangle the reaction to anticipated and unanticipated cost changes, we leverage data on wholesale gasoline futures traded on the New York Mercantile Exchange (NYMEX). The presence of a futures market allows us to project expectations of future wholesale costs for the firms in our market.

To make these projections, we assume that firms are engaging in regression-like predictions of future wholesale costs, and we choose the 30-day-ahead cost as our benchmark.[29] Using station-specific wholesale costs, we regress the 30-day lead wholesale cost on the current

[28] We have also estimated specifications that add lagged prices. Consistent with past prices and demand shocks affecting current choices, lagged prices are significant and the lagged share coefficient is similar.

[29] Futures are specified in terms of first-of-the-month delivery dates. To convert these to 30-day-ahead prices, we use the average between the two futures, weighted by the relative number of days to the delivery date.

wholesale cost and the 30-day-ahead future. In particular, we estimate the following equation.

$$c_{nt+30} = \alpha_1 c_{nt} + \alpha_2 F_t^{30} + \gamma_n + \epsilon_{nt} \tag{B.1}$$

Here, c_{nt+30} is the 30-day-ahead wholesale cost for station n, F_t^{30} is the 30-day-ahead forward contract price at date t, and γ_n is a station fixed effect. We use the estimated parameters to construct expected 30-day-ahead costs for all firms: $\hat{c}_{nt+30} = \hat{\alpha}_1 c_{nt} + \hat{\alpha}_2 F_t^{30} + \hat{\gamma}_n$. The unexpected cost, or cost shock, is the residual: $\tilde{c}_{nt+30} = c_{nt+30} - \hat{c}_{nt+30}$.

For robustness, we construct a number of alternative estimates of expected costs, including a specification that makes use of all four available futures. However, we found that these alternative specifications were subject to overfitting; the estimates performed substantially worse out-of-sample when we ran the regression on a subset of the data. Our chosen specification is remarkably stable, with a mean absolute difference of one percent when we use only the first half of the panel to estimate the model. Expected costs constitute 74.6 percent of the variation in costs (R^2) in our two-year sample, which includes a large decline in wholesale costs due to several supply shocks in 2014.

In Section 3.3.2, we consider only the simple cut between unexpected and expected costs to focus attention on this previously unexplored dimension of pass-through. In retail gasoline markets, costs are highly correlated, with common costs tending to dominate idiosyncratic costs at moderate frequencies. For robustness, we have estimated the cost pass-through (i) using only common costs and (ii) controlling for the mean cost of rival brands (in the same county). In either scenario, we find estimates that are very similar.

A Note on 30-Day-Ahead Expectations

One of the challenges in discussing expectations is that they change each day with new information. News about a cost shock 30 days from now may arrive at any time within the next 30 days, if it has not arrived already. Therefore, any discussion of an "unexpected" cost shock must always be qualified with an "as of when." Given previous findings in the gasoline literature indicating that prices take approximately four weeks to adjust, a 30-day-ahead window seems an appropriate one to capture most of any anticipatory pricing behavior. Additionally, our findings support this window as reasonable in this context. We see no relationship between unexpected or expected costs and the price 30 days prior.[30]

[30] We interpret slight deviations from zero as arising from an underlying correlation in unobserved cost shocks.

C Estimation Details

C.1 Reducing the Computational Burden

Though the distribution of unobserved choices is identified, solving for the pattern of choices in estimation is another matter. The traditional approach is to "concentrate out" the distribution of unobserved heterogeneity while using a contraction mapping to solve (implicitly) for the shares of the type-0 consumers (as in Berry et al., 1995). In our setting, the assumption of single-product affiliation allows us to reduce the computational burden, as the full distribution of choice patterns in each market can be calculated directly after solving a system of equations in two variables. Thus, we reduce the number of unknowns in each market from J to 2. This may be used to speed up estimation by implementing a non-linear equation solver or a (modified) contraction mapping.

In Section 4.1, we showed that the choice patterns can be expressed in terms of the $J+1$ parameters $\{s_{0t}(j)\}$ in each market. We now show that the system reduces to two parameters in each market, where the remaining $J-1$ parameters are solved for by a quadratic function.

Under the assumption of single-product affiliation ($\sigma_{jt}(z) = 0 \, \forall \, z \neq j$), we obtain

$$\sum_z r_{zt} \cdot s_{0t}(z) \exp\left(\sigma_{jt}(z)\right) = \sum_z r_{zt} s_{0t}(z) + \left(\exp\left(\sigma_{jt}(j)\right) - 1\right) r_{jt} s_{0t}(j). \tag{C.1}$$

By substituting this expression into equation (5), along with expressions (10) and (11), we can obtain a quadratic equation for each of the $\{s_{0t}(j)\}$:

$$\begin{aligned} 0 = \lambda \frac{1}{s_{0t}(0)} r_{jt} \Bigg[& \left(\exp\left(\sigma_{jt}(j)\right) - 1\right) \boldsymbol{s_{0t}(j)^2} \\ & + \left(\left(\exp\left(\sigma_{jt}(j)\right) - 1\right) \left(S_{jt} - \lambda r_{jt}\right) + \lambda \frac{1}{s_{0t}(0)} \sum_{z \in \{0\} \cup J} r_{zt} s_{0t}(z) + (1-\lambda) \right) \boldsymbol{s_{0t}(j)} \\ & - \lambda \sum_{z \in \{0\} \cup J} r_{zt} s_{0t}(z) - (1-\lambda) s_{0t}(0) \Bigg]. \end{aligned}$$

Conditional on the dynamic parameters and observables, there are only two remaining unknowns: $s_{0t}(0)$ and $\sum_{z \in \{0\} \cup J} r_{zt} s_{0t}(z)$. Thus, we can solve for $\{s_{0t}(j)\}$ in each market using the quadratic formula. As $\{\delta_{jt}\}$ are identified conditional on these choice probabilities, we can obtain these mean utility parameters by solving for only two unknowns in each market, regardless of the number of products.

Table C.1: Data-Generating Process

Parameter/Variable	Value	Description
β	2	Demand intercept
α	-3	Mean price coefficient
ξ_j	(2, 2, 1, 1, 0, 0)	Product fixed effects
$\Delta\xi_{jrt}$	U(0,1)	Error term
D_i	N(0,1)	Demographic variable
mc_{jrt}	U(1,3)	Marginal costs
p_{jrt}	mc_{jrt} + U(0,2)	Prices
z_{jrt}	mc_{jrt} + N(0,0.01)	Instrument for price

C.2 Identification: Monte Carlo Exercise

We use Monte Carlo simulations to demonstrate that (1) our model does not incorrectly attribute persistent between-consumer heterogeneity to state dependence, and (2) the model recovers the correct parameters when state-dependent demand is present. Therefore, our method captures the presence of dynamic demand even when persistent unobserved heterogeneity is generated from more complex preferences than we model. As discussed in the paper, this is important because dynamic demand introduces bias not only through the demand elasticities but also through the firms' first-order conditions, which account for future demand. Further, the simulations help illustrate how state dependence can generate biased elasticities when not included in the model.

For the data-generating process, we assume that the indirect utility consumer i receives from purchasing product $j > 0$ in region r and period t is:

$$u_{ijrt} = (\beta + \pi_1 D_i) + (\alpha + \pi_2 D_i)p_{jrt} + \xi_j + \Delta\xi_{jrt} + \epsilon_{ijrt}, \tag{C.2}$$

where p_{jrt} is the retail price, ξ_j denotes product fixed effects, $\Delta\xi_{jrt}$ is a structural error term, and ϵ_{ijrt} is a consumer-specific logit error term.[31] A consumer that selects the outside good receives $u_{i0rt} = \epsilon_{i0rt}$.

Persistent consumer-specific preferences are governed by the parameters (π_1, π_2) and load onto the demographic variable D_i (e.g., "income"). This variable captures heterogeneity in consumer preferences regarding the propensity to buy any good (π_1) and price sensitivity (π_2).

We consider five different choices of the parameters (π_1, π_2), including the baseline standard logit. Specifically, for (π_1, π_2), we choose $(0, 0)$, $(4, 1)$, $(8, 0)$, $(0, 2)$, and $(8, 2)$. These specifications help illustrate how different types of persistent between-consumer heterogeneity might introduce bias into our model. We assume that D_i is distributed according to the normal distribution $N(0, 1)$. For each specification, we simulate 50 regions and 100 periods, with 6 products in each region-period. We draw 500 individuals in each region and use the standard choice

[31] The product fixed effects allow for persistent (shared) preferences for specific products, but these (as might be expected) are easy to control for.

Table C.2: Monte Carlo Results: Dynamic Model Estimates

	True Parameters			Elasticity	Estimate
Specification	π_1	π_2	λ	ϵ_D	$\hat{\lambda}$
1	0	0	0	-8.94	0.000
2	4	1	0	-4.29	0.000
3	8	0	0	-8.54	0.000
4	0	2	0	-1.78	0.000
5	8	2	0	-1.86	0.000
6	0	0	0.1	-5.64	0.095
7	0	0	0.3	-4.45	0.340
8	0	0	0.5	-3.85	0.496
9	0	0	0.7	-3.41	0.686
10	0	0	0.9	-3.05	0.888

probability equations to calculate shares. We choose parameter values that result in reasonable mean shares across specifications, ranging from 0.035 to 0.092. Since, for the purposes of this exercise, we are interested in estimating demand, we simulate prices by randomly drawing markups and adding them to marginal costs. The parameter values used for the simulations are reported in Table C.1.

In addition to these static specifications, we also choose five dynamic specifications where $(\pi_1, \pi_2) = (0, 0)$. As in the paper, we assume that a fraction of consumers, λ, receive a preference shock, $\overline{\xi}$, for the product they purchased in the previous period. The remaining $1 - \lambda$ consumers choose according to the standard logit. We consider $\lambda \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ and set $\overline{\xi} = 8$. Thus, aside from the parameters π_1, π_2, and λ, all ten specifications share identical parameterizations.

Table C.2 reports the 10 specifications in our simulations. To illustrate the magnitudes of the impact of consumer heterogeneity and state dependence, we report in this table the true median own-price elasticity. Because the richer preferences shift the marginal consumers, specifications 2 through 10 have less elastic demand than the baseline logit.

We estimate the model following the approach detailed in the body of the paper, assuming the structure of the dynamic demand model. We use the same estimation code for each specification, including the same (nonzero) initial parameter values, and we confirm that we reach the minimum. The last column of Table C.2 reports the estimated value for $\hat{\lambda}$, the share of consumers who are affiliated. The model recovers the true share of affiliated consumers, even in the presence of random coefficients (which are not directly modeled in our empirical specification). In specifications 1 through 5, the model estimates that 0.0 percent of customers are subject to state dependence. By contrast, the estimates in specifications 6 through 10 are close to the true values, with small differences that can be explained by sampling error and our modest sample size.[32]

[32]The estimated values for $\overline{\xi}$ for specifications 6 through 10 are 8.0, 8.2, 7.9, 8.3, and 8.3, close to the true value of 8. To ensure that the state dependence shock has bite (i.e., if $\overline{\xi} = 0$, λ is irrelevant), we set the minimum value of $\overline{\xi}$ to 3 in estimation. Specifications 1–5 are above this lower bound, but it is irrelevant as λ is estimated to be zero.

Table C.3: Unobserved Heterogeneity and Bias in Static Logit Estimates

	Random Coefs.		Dynamics	Price Coefficient			Demand Elasticity		
Specification	π_1	π_2	λ	α	$\hat{\alpha}$	Std. Err.	ϵ_D	$\hat{\epsilon}_D$	% Bias
1	0	0	0	-3	-3.000	0.003	-8.94	-8.95	0.0
2	4	1	0	-3	-1.542	0.015	-4.29	-4.44	3.5
3	8	0	0	-3	-2.640	0.016	-8.54	-7.71	-9.7
4	0	2	0	-3	-0.679	0.014	-1.78	-1.93	8.4
5	8	2	0	-3	-0.688	0.014	-1.86	-1.94	4.7
6	0	0	0.1	-3	-2.099	0.008	-5.64	-6.21	10.0
7	0	0	0.3	-3	-1.730	0.010	-4.45	-5.02	12.7
8	0	0	0.5	-3	-1.554	0.011	-3.85	-4.43	15.3
9	0	0	0.7	-3	-1.436	0.012	-3.41	-4.04	18.4
10	0	0	0.9	-3	-1.348	0.013	-3.05	-3.73	22.2

Notes: Table C.3 reports the results when the DGP has random coefficients (2)–(5) and has dynamic demand (6)–(10). The true price coefficient is α, and the estimated coefficient from the logit regression is $\hat{\alpha}$. We estimate the (misspecified) regression equation $\ln(s_{jrt}/s_{0rt}) = \beta + \alpha p_{jrt} + \xi_j + \Delta\xi_{jrt}$. We instrument for price with marginal costs plus a small error term (see Table C.1 for details on z and mc). We also report, in the last three columns, the true own-price demand elasticity (ϵ_D), the estimated own-price demand elasticity ($\hat{\epsilon}_D$), and the percent bias of the estimate. The results show that both random coefficients and state dependence generate bias in static logit estimates.

As an additional way to illustrate how random coefficients can introduce bias in (misspecified) demand estimates, we provide the estimated price coefficients and demand elasticities from a static logit model in Table C.3. As expected, specifications 2 through 5 have substantial bias in the estimated price coefficient (the structural parameter), but the specifications have somewhat more modest bias in own-price elasticities. By contrast, the specifications with (unaccounted-for) dynamic demand yield greater bias in the own-price elasticities, even for small values of λ.

We conclude three things from these simulations. First, our estimation strategy is capable of correctly recovering the true degree of state dependence when the model is correctly specified. Second, our dynamic model does not attribute persistent between-consumer heterogeneity to state dependence in demand. Third, state dependence in demand is as capable of generating bias when omitted from the model as persistent unobserved heterogeneity (and perhaps more so, when accounting for the firms' first-order conditions). In many cases, within-consumer state dependence may be the most relevant feature of the economic environment for the question at hand. In cases where persistent between-consumer unobserved consumer heterogeneity is also important, our approach may be used as a test for the presence of dynamic demand, and the results can inform the best model to use for a more detailed analysis.

D Empirical Application: Additional Results

D.1 Additional Tables and Figures

Here we include time-series plots of total market shares, prices, and the instrument (Figure D.1), summary statistics at the observation level (Table D.1), and summary statistics by brand (Table D.2).

Figure D.1: Shares and Prices

(a) Total Market Shares and Prices

(b) Instrument and Prices

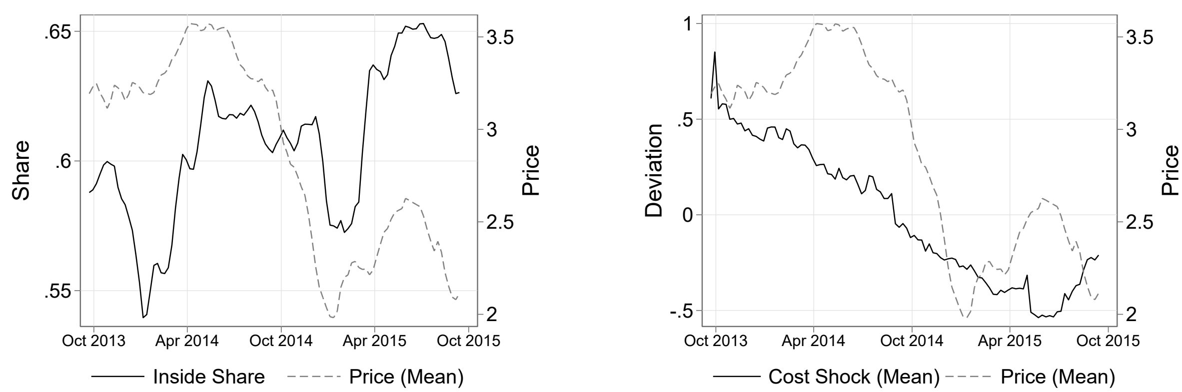

Notes: Panel (a) displays the sum of market shares for all brands, excluding only the outside option, plotted with the average price over the period in our sample. Both lines indicate seasonality, with peaks occurring during the summer. Panel (b) plots the constructed instrument against the average price in our sample. Overall, there is a strong negative correlation between the instrument and average prices.

Table D.1: Retail Gasoline in Kentucky and Virginia: Oct 2013 – Sep 2015

Statistic	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max	N
Share	0.137	0.103	0.000	0.060	0.183	0.688	110,844
Price	2.871	0.529	1.715	2.384	3.311	4.085	110,844
Wholesale Price	2.257	0.527	1.245	1.754	2.673	3.545	110,844
Wholesale FE	2.261	0.031	2.206	2.230	2.293	2.366	110,844
Margin ($/gal)	0.206	0.114	-0.440	0.132	0.272	1.048	110,844
Num. Stations	5.050	6.780	1.000	2.000	6.000	79.000	110,844

Notes: Table provides summary statistics for the observation-level data in the analysis. Prices and margins are reported in dollars per gallon. The greatest number of stations a brand has in a single county in our data is 79. The 25th percentile is 2, and we have several observations of a brand with only a single station in a market. The variable Wholesale FE is the average wholesale price for a brand within a county. We interact this variable with the U.S. oil production data to generate an instrument for price in the demand estimation. Negative margins occur in 2.7 percent of observations.

Table D.2: Summary of Brands

Brand	Cond. Share	Share	Num. Markets	Num. Stations	Margins ($/gal)
Marathon	0.179	0.099	134	5.3	0.21
Sheetz	0.176	0.027	37	1.7	0.17
Speedway	0.174	0.028	39	3.7	0.18
Wawa	0.159	0.014	22	3.2	0.12
Exxon	0.153	0.073	116	4.6	0.25
7-Eleven	0.134	0.023	41	6.8	0.18
Shell	0.133	0.089	163	4.2	0.22
FRINGE	0.132	0.119	233	8.7	0.19
Pilot	0.124	0.010	21	1.4	0.13
BP	0.116	0.059	124	3.3	0.21
Loves	0.112	0.007	15	1.0	0.20
Valero	0.109	0.027	59	3.5	0.21
Thorntons	0.106	0.004	9	5.9	0.14
Huck's	0.106	0.004	10	1.9	0.15
Sunoco	0.090	0.012	32	4.7	0.28
Citgo	0.074	0.010	34	4.0	0.25

Notes: Table provides summary statistics by brand. The FRINGE brand is a synthetic brand created by aggregating brand-market observations that fail the brand-selection screen. A brand is retained as a standalone brand if it appears in at least ten markets and either has an average share of at least 2 percent across markets or an average share of at least 10 percent conditional on being present. Even for retained brands, we assign a brand to FRINGE within a market if its average share in that market is below 5 percent. Unbranded observations are also assigned to FRINGE. We additionally make seven market-specific relabelings where the price series and share series appear mismatched.

D.2 Details on Forward Simulations

Figure D.2: Simulated and Fitted Values

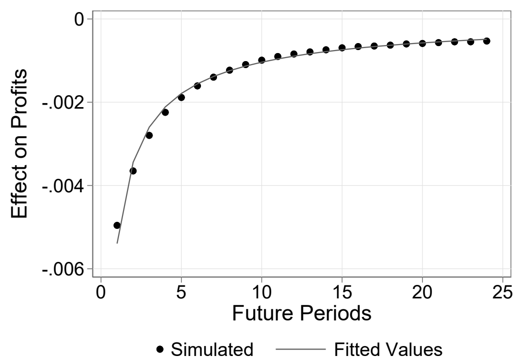

Notes: Figure displays the equilibrium effect of a marginal increase in price on future profits. An increase in price decreases future profits. This effect is diminishing over time. The points display simulated values from the data, and the line displays fitted values from a regression model.

To test whether firm expectations are consistent with expected profits under our supply-side approach, we simulate equilibrium prices and profits using marginal deviations for individual prices. We perturb the price of a specific brand in a given market-week by lowering that brand's price by \$0.01. We recompute shares in that period, then we compute equilibrium prices over the next 24 weeks. We repeat this experiment for all brands, markets, and periods in the sample used for the counterfactual. We limit our analysis to the 75 markets and 51 weeks affected by our merger counterfactual.

We use the difference in simulated profits and realized profits to compute the marginal effect of a \$0.01 change in price on profits τ periods in the future. By scaling the difference by the magnitude of the price change, we obtain an estimate of the effect on current-period profits and each period-specific component of the derivative $\frac{\partial \pi_{j(t+\tau)}}{\partial p_{jt}}$. The simulated effect on current-period profits from a price change matches the static derivative of the model. To calculate the discounted future values, we use a weekly discount rate of $\beta = 0.999$, which corresponds to an annual discount rate of 0.949. We use the discounted values of $\frac{\partial \pi_{j(t+\tau)}}{\partial p_{jt}}$ to obtain an estimate of the full value of $\beta \frac{\partial E[V_j(\cdot)|\cdot]}{\partial p_{jt}}$.

The dynamic equilibrium simulation shows that an increase in price today has a negative effect on future profits, and this effect shrinks over time. To calculate the net present value of the effect on profits, we use a fitted model to predict effects beyond 24 weeks and to mitigate the effect of simulation noise. Specifically, we fit a curve of the form $\ln\left(-\ln\left(E_t[\frac{\partial \pi_{j(t+\tau)}}{\partial p_{jt}}]\right)\right) = \omega_0 + \omega_1 \ln(t) + \epsilon_t$. We estimate $\omega_0 = 1.653$ and $\omega_1 = 0.119$ with least squares regression, using the 24 mean values for each period. Figure D.2 plots the simulated effects on future profits against the fitted values. The fitted values explain 99 percent of the variation in the mean values.

D.3 Counterfactuals: Isolating the Price Effects of Consumer Inertia

In the main text, we use merger simulations to quantify horizontal market power in the presence of consumer inertia. Here, we isolate dynamic market power, that is, the pricing effects of consumer inertia. To do so, we employ counterfactuals where we change the strength of affiliation, $\overline{\xi}$, and the share of consumers who are subject to state dependence, λ. These simulations help contextualize the economic significance of inertia and the relative effects we find from mergers.

Specifically, we simulate one counterfactual in which $\overline{\xi}$ increases by 10 percent and a second in which λ_m increases by 0.10 in every market. The latter corresponds to an average change in λ_m of 16 percent. Thus, we consider the impact of both the strength of affiliation and the share of consumers subject to inertia on equilibrium prices. We use the same 75 markets as in Section 5 to facilitate a comparison to the horizontal market power effects in Table 7.

Table D.3 reports the results from the counterfactual scenarios. The first three columns report the equilibrium effects of an increase in the strength of state dependence, and the next three columns report the effects of an increase in its prevalence. The results are reported by brand, and they are sorted by the mean share in the 75 relevant markets from the previous subsection. Thus, Marathon and BP appear as the two largest brands in these markets. If the state-dependent utility shock increased by 10 percent, then overall prices would increase by 4.7 percent and average shares would fall by 1.8 percent. All firms realize an increase in profits, with some smaller brands (Huck's and Thorntons) realizing the lowest increases in prices. There is also variation in the change in shares, with Pilot realizing an increase in shares (and the largest gain in profits), while Sheetz experiences no change in equilibrium.

Increasing λ yields smaller effects on prices (1.7 percent), though firms increase shares by 6.1 percent and profits still increase by a meaningful amount. Interestingly, the effects on prices, shares, and profits are not perfectly correlated across the two counterfactuals. For example, Sheetz has the second-largest increase in profits from a change in the strength of affiliation but the third-smallest profit change from an increase in the prevalence of state dependence. These results suggest that multiple dimensions of a firm's relative position in the market influence the impact of consumer inertia.

In terms of magnitudes, these effects on prices and profits are in the same ballpark as the effects of mergers between Marathon and BP (Table 7). In rough terms, increasing the strength of affiliation has a similar effect on prices as the joint pricing merger, though shares do not fall by as much and profits, accordingly, increase by more. By contrast, increasing the share of consumers affected by inertia by 0.10 has effects on prices, shares, and profits that are roughly three-fifths of those of a brand consolidation merger. Thus, the model allows us to quantify the relative effects of dynamic incentives and horizontal competition when evaluating market power. Our findings suggest that policies that produce even modest changes in the prevalence or magnitude of switching costs (or consumer inertia, more generally) may have a similar impact

Table D.3: Equilibrium Effects of State Dependence

	Increase $\bar{\xi}$ by 10 Percent			Increase λ by 0.10		
Brand	Price	Share	Profit	Price	Share	Profit
Marathon	4.51	-2.4	45.8	1.90	6.5	29.0
BP	4.58	-1.8	52.2	1.91	4.9	27.7
FRINGE	4.50	-0.6	63.3	1.42	6.3	29.1
Shell	4.90	-2.0	53.6	1.72	5.6	26.5
Speedway	5.04	-2.0	60.6	1.67	7.3	29.7
Exxon	5.51	-3.3	48.8	1.66	7.0	24.9
Valero	4.40	-1.9	51.5	1.77	3.0	26.7
Sheetz	5.71	0.0	78.7	1.03	10.1	26.3
Loves	4.52	-1.3	51.1	2.14	5.4	31.4
Pilot	4.42	2.2	126.1	1.04	8.9	41.5
Huck's	3.82	-2.4	67.3	1.56	5.7	36.6
Sunoco	5.11	-2.0	42.4	2.73	-3.8	19.8
Citgo	5.67	-1.0	59.3	1.36	5.9	21.5
Thorntons	4.00	-1.9	56.2	1.76	6.5	32.6
Overall	4.68	-1.8	53.6	1.73	6.1	28.1

Notes: Table displays the mean percent changes in prices, shares, and profits from counterfactual scenarios with different portions of consumers affected by state dependence. The first three columns report the equilibrium effects with a greater strength of state dependence, and the next three columns report the effects with a greater share of consumers affected by dependence.

on welfare as large changes to market structure, such as those that occur through mergers.

E Theoretical and Numerical Analysis: Additional Details

E.1 Numerical Simulation Methodology

The number of unknowns in the system is $J + J + J \times J$, for p, r, and $\frac{dp}{dr}$. The law of motion in the steady state gives us J restrictions ($r = f(p, r)$). This allows us to solve for r given p. We solve for p and $\frac{dp}{dr}$ using steady-state conditions.

We implement the following procedure to solve numerically for the steady state:

1. Provide an initial guess for the matrix $\frac{dp}{dr}$.

2. Solve for steady-state values of p, r, and $\frac{dV_k(r)}{dr}$ given the guess for $\frac{dp}{dr}$. Use the J restrictions implied by the first-order conditions (one for each product j)

$$\frac{dV_k(r')}{dr'} \cdot \frac{dr'}{dp_j} = -\frac{1}{\beta} \sum_{l \in J_k} \frac{\partial \pi_l}{\partial p_j},$$

   to solve for $\frac{dV_k(r)}{dr}$, where $\frac{dV_k(r)}{dr} = \frac{dV_k(r')}{dr'}$ in the steady state. Note that π_k, in this notation, is equal to the sum of profits from all products by a firm.

3. Take the numerical derivative of p with respect to r. Approximate the numerical derivative by slightly perturbing r: $\tilde{r}_j = r + \epsilon_j$, where j indicates a perturbation in the j^{th} element. Re-solve for p using the first-order condition. Calculate

$$\frac{dp}{dr_j} \approx \frac{p^*(r + \epsilon_j) - p^*(r - \epsilon_j)}{2\,|\epsilon_j|}$$

   Stack these vectors horizontally to obtain an approximation for $\frac{dp}{dr}$.

4. Calculate the absolute distance between the approximation of $\frac{dp}{dr}$ calculated in the previous step and the initial guess for $\frac{dp}{dr}$.

5. If the calculated distance falls below a critical value, then the solution is found. If not, update the guess for $\frac{dp}{dr}$ as the average between the initial guess and the approximation calculated in step 3. Repeat steps 1-4 above until a solution is found.

E.2 Parameter Values

We attempt to simulate data from the support of parameters that produce reasonable outcomes for markups and shares. We employ a "shotgun" approach, generating simulations with many different parameters and selecting only the markets that meet certain criteria. We first take Halton draws of the demand parameters such that $\xi \in [0, 10]$, $\overline{\xi} \in [0, 10]$, $\alpha \in [-10, 0]$, and set each firm's marginal cost to one. For each draw of these demand parameters, we construct

Table E.1: Simulation Parameter Summary Statistics

	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max	N
λ	0.38	0.20	0.05	0.20	0.55	0.70	6566
α	-5.21	2.56	-9.96	-7.31	-3.12	-0.44	6566
ξ	5.27	2.60	0.04	3.29	7.37	9.99	6566
$\bar{\xi}$	2.92	1.49	0.00	1.76	4.14	5.97	6566

Notes: Table displays summary statistics for demand parameters for the 6,566 markets used in the numerical simulations. These markets were generated from a broader set of parameter values and selected if the resulting three-firm markets had firm shares between 0.05 and 0.30 (yielding an outside share between 0.10 and 0.85), markups between 0.05 and 0.75, and converged for all values of $\lambda \in \{0.05, 0.1, 0.15, \ldots, 0.70\}$. See the text for additional details.

three-firm markets for $\lambda \in \{0.05, 0.1, 0.15, \ldots, 0.70\}$. We then restrict the analysis to markets where firms have shares between 0.05 and 0.30 (yielding an outside share between 0.10 and 0.85) and Lerner index markups ($\frac{p-c}{p}$) between 0.05 and 0.75.[33] Finally, to avoid composition effects, we only analyze markets with demand parameters that converged for all values of λ. We set the upper bound of λ to be 0.70, because for higher values of λ there were often markets where the post-merger equilibrium did not converge or pre-merger markups fell above 0.90. The data-generating process yields 6,566 markets whose parameters are summarized in Appendix Table E.1.

E.3 Summary of Markets: Dynamic and Horizontal Market Power

Affiliation in our model can affect equilibrium prices through its influence on *dynamic market power*—the extent to which a firm sets higher prices due to the presence of state-dependent consumers—and through its influence on *horizontal market power*—the extent to which competitors constrain prices.

Figure E.1 plots the effects of affiliation and reduced competition on prices. The plots employ simulation results from the 6,566 markets, or 469 baseline parameter values of $\xi, \bar{\xi}$, and α that converged for all $\lambda \in [0.05, 0.70]$. Panel (a) measures the pricing effects of consumer inertia by plotting price effects as a function of λ. Percent changes are calculated relative to no consumer affiliation ($\lambda = 0$) while holding fixed the other parameters in the model. On average, prices increase with the fraction of customers prone to affiliation for values of $\lambda \leq 0.55$ and decrease thereafter. In our simulations, both market share and profits increase monotonically with λ (not shown). On average, affiliation results in moderate price increases compared to a static demand model. However, the impact of dynamic market power can be quite substantial depending on the underlying demand parameters. The 90–10 range of outcomes is plotted with

[33] The range for each parameter is selected such that parameter values just outside the bounds of the range result in outcomes that often fall above or below our share and markup criteria.

Figure E.1: Potential Price Effects

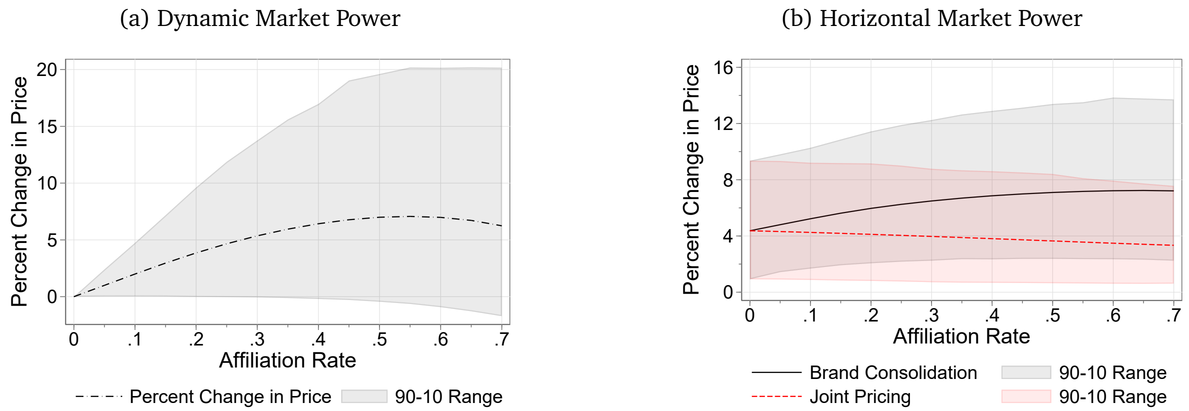

Notes: Panel (a) displays the mean percent price increase for a three-firm oligopoly relative to the baseline model with no dynamics in consumption ($\lambda = 0$). Panel (b) displays the mean percent price increase of a merger to a duopoly for two types of mergers: joint pricing control and brand consolidation, for different values of λ. The plots reflect 469 baseline parameter values of $(\xi, \bar{\xi}, \alpha)$ that converged for all $\lambda \in \{0.05, 0.1, 0.15, ..., 0.70\}$, or 6,566 markets in total.

the transparent area. When $\lambda > 0.4$, the 90th percentile market exhibits prices more than 15 percent higher than in the static benchmark.

In panel (b), we plot the effect of horizontal market power on prices across different values of λ. We measure horizontal market power by comparing the prices in the symmetric three-firm oligopoly to the post-merger prices in both joint pricing and brand consolidation mergers. In our symmetric setting, brand consolidation mergers provide greater horizontal market power relative to joint pricing mergers across all levels of λ. There is substantial overlap in the 90–10 percentile range across the two merger types, demonstrating that both can enable comparable levels of horizontal market power. Panel (b) also shows that horizontal market power slightly decreases with λ, on average, for joint pricing mergers, while it increases in brand consolidation mergers. This highlights that the interaction between consumer inertia and a decrease in competition depends critically on the mechanism through which competition is reduced.

Table E.2 provides additional statistics about the horizontal mergers performed in each market. The simulations demonstrate that joint pricing and brand consolidation mergers can have diverging effects in the presence of consumer inertia. The average percentage merger price effect is 3.8 percent for joint pricing mergers and 6.5 percent for brand consolidation mergers. The table presents the distribution of elasticities; the mean dynamic elasticity is -3.86, while the misspecified static elasticity is -4.7.

Table E.2: Simulation Summary Statistics

	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max	N
Pre-Merger Price	1.43	0.42	1.10	1.19	1.50	3.97	6566
Pre-Merger Markup	0.26	0.14	0.09	0.16	0.33	0.75	6566
Pre-Merger Market Share	0.17	0.07	0.05	0.11	0.24	0.30	6566
HHI: Pre-Merger	1067.35	776.19	76.28	339.49	1726.75	2695.59	6566
Δ HHI	711.57	517.46	50.85	226.33	1151.17	1797.06	6566
Joint Pricing Merger Δ	3.84	3.37	0.25	1.30	5.32	20.57	6566
Joint Pricing Non-Merging Price Δ	0.91	1.17	0.00	0.11	1.26	7.86	6566
Brand Consolidation Price Δ	6.49	4.29	0.23	3.43	8.55	25.48	6566
BC Non-Merging Price Δ	0.81	1.11	-1.14	0.07	1.17	8.64	6566
Prediction Bias: Joint Pricing	1.47	1.70	0.00	0.35	1.97	13.49	6566
Prediction Bias (pctg.): Joint Pricing	67.34	99.87	0.10	12.54	74.27	792.05	6566
Prediction Bias: Brand Consolidation	-1.17	1.46	-8.63	-1.77	-0.22	4.34	6566
Prediction Bias (pctg.): Brand Consolidation	-19.70	24.00	-76.01	-35.93	-4.16	181.04	6566
Dynamic Elasticity: Unaffiliated	-5.76	2.37	-11.96	-7.55	-3.77	-1.29	6566
Dynamic Elasticity: Affiliated	-2.13	1.90	-9.74	-3.14	-0.74	-0.08	6566
Dynamic Elasticity: Weighted	-3.86	1.98	-10.03	-5.15	-2.24	-0.59	6566
Static Elasticity	-4.74	2.08	-10.71	-6.15	-3.00	-1.34	6566

Notes: Markup is defined as $\frac{p-c}{p}$. Δ HHI is calculated at the pre-merger shares. Merger Price Δ is the percentage price increase from the merger. Prediction Bias is the static prediction minus the dynamic prediction, in percentage points. Prediction Bias (pctg.) is the Prediction Bias divided by the dynamic Merger Price Δ. The weighted dynamic elasticity is the average of the unaffiliated and affiliated elasticities weighted by the fraction of the firm's customers of each type.

E.4 Example Markets

To illustrate the impact that affiliation can have on pricing incentives, we plot the equilibrium prices for two different sets of utility parameters in Figure E.2. Panel (a) plots the equilibrium prices for a monopolist and a three-firm symmetric market, for increasingly large values of λ and otherwise identical demand parameters. The equilibrium prices increase with the rate of affiliation for values of λ less than 0.5 and then decrease again. This figure highlights the potential importance of affiliation for equilibrium prices and shows that the effect on price may be non-monotonic in the proportion of consumers prone to affiliation. The non-monotonicity is the result of two countervailing effects. At low levels of λ, firms face increasingly inelastic demand and therefore increase prices. As λ increases past 0.5, however, the incentive to invest in future demand swamps the elasticity effect and puts downward pressure on prices. Note that firms' profits continue to increase, even as prices begin to decrease.

Panel (b) shows that the investment incentive can dominate and grow stronger for all values of λ. Furthermore, the investment incentive can be blunted by competition. Prices remain relatively flat for all values of λ in the three-firm market, whereas in the monopoly market the investment incentive is stronger and prices decrease at a faster rate. Thus, the relationship between equilibrium prices and dynamics may depend upon market structure.

Figure E.2: Monopoly and Oligopoly Prices with Consumer Affiliation

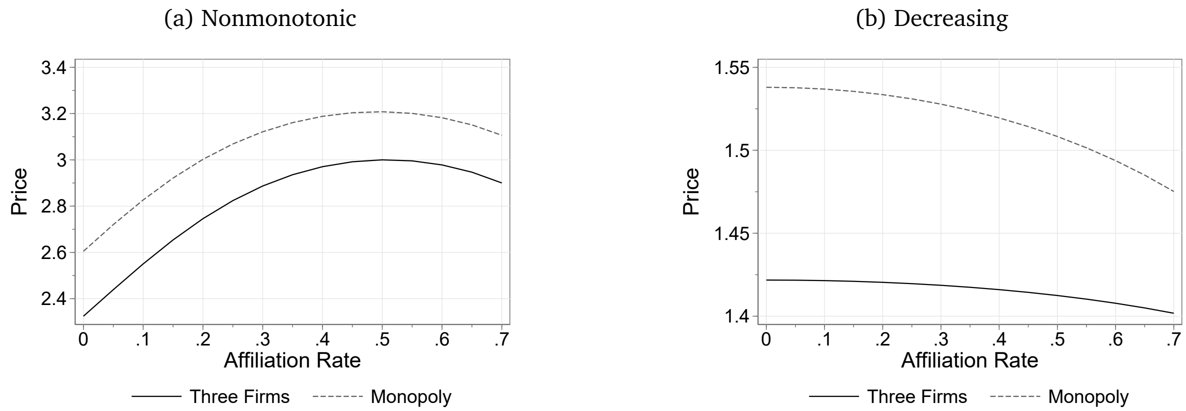

Notes: Panel (a) is generated using the following parameter values: $\xi = 0.04$, $\overline{\xi} = 4.15$, and $\alpha = -0.84$. Panel (b) is generated using $\xi = 2.33$, $\overline{\xi} = 0.04$, and $\alpha = -2.73$. Marginal cost is set to one for both figures.

E.5 Ambiguous Effects of Affiliation: Analytical Results for Monopoly

Here, we present analytical results showing how affiliation can have ambiguous effects on prices. We analyze steady-state prices in a monopoly market (with an outside good) to show how habit-forming consumers affect optimal prices and markups. The monopolist sells product j and the outside good is indexed as product 0. In the steady state, $r_{jt} = r_{j(t+1)} = r_j$ and $c_{jt} = c_{j(t+1)} = c$. Given the model detailed in the body of the paper, the monopolist's steady-state number of affiliated consumers, r_j^{ss}, is:

$$\begin{aligned} r_j^{ss} &= r_0^{ss} s_j(0) + r_j^{ss} s_j(j) \\ \implies r_j^{ss} &= \frac{r_0^{ss} s_j(0)}{1 - s_j(j)}. \end{aligned}$$

In the monopoly case, we also have that $r_0^{ss} = 1 - r_j^{ss}$. It follows that we can express r_j^{ss} as a function of the choice probabilities for product j.

$$r_j^{ss} = \frac{s_j(0)}{1 - s_j(j) + s_j(0)} \tag{E.1}$$

Using the steady-state value of affiliated consumers, the fraction of consumers subject to affiliation, λ, and the aggregate share equation, S_j we can solve for the steady-state pricing function.

The steady-state period value is:

$$\begin{aligned} V^{ss}(r^{ss}, c^{ss}) &= (p^{ss} - c^{ss})((1 - \lambda + \lambda r_0^{ss})s_j(0) + \lambda r_j^{ss} s_j(j)) + \beta V^{ss} \\ &= (p^{ss} - c^{ss})(s_j(0) + \lambda r_j^{ss}(s_j(j) - s_j(0))) + \beta V^{ss} \\ &= \frac{p^{ss} - c^{ss}}{1 - \beta} \cdot (s_j(0) + \lambda r_j^{ss}(s_j(j) - s_j(0))). \end{aligned}$$

This equation represents the monopolist's discounted profits, conditional on costs remaining at its current level. Thus, profits are increasing in both λ and the difference in choice probabilities of affiliated and unaffiliated consumers. These results are straightforward: affiliated consumers are profitable. Also, note that a model with no affiliation is embedded in this formulation ($\lambda = 0$ and $s_j(j) = s_j(0)$), in which case profits are simply the per-unit discounted profits multiplied by the firm's market share. Because the steady-state can be expressed entirely in terms of product j choice probabilities and affiliated customers, we simplify the following notation: $s_j(j) = s^j$, $s_j(0) = s^0$, and $r_j = r$.

Maximizing the steady-state value with respect to p^{ss} yields the firm's optimal pricing function:

$$p^{ss} = c^{ss} + \underbrace{\frac{-s^0 - \lambda r^{ss}\left(s^j - s^0\right)}{\frac{ds^0}{dp} + \lambda \frac{dr^{ss}}{dp}\left(s^j - s^0\right) + \lambda r^{ss}\left(\frac{ds^j}{dp} - \frac{ds^0}{dp}\right)}}_{m\text{ = markup of price over marginal cost}}. \tag{E.2}$$

The second term, m, on the right-hand side of equation (E.2) captures the extent to which the firm prices above marginal cost (in equilibrium). As this markup term depends upon choice probabilities, it is implicitly a function of price. Thus, as in the standard logit model, we cannot derive an analytical solution for the steady-state price. Nonetheless, we derive a condition below to see how markups are impacted by consumer affiliation. In the usual case, m will be declining in p, ensuring a unique equilibrium in prices.

Are markups higher or lower in the presence of affiliation? When affiliation is absent, $\lambda = 0$ and $s^j = s^0$, equation (E.2) reduces to the first-order condition of the static model, $p^{ss} = c^{ss} - \frac{s^0}{ds^0/dp}$. Denoting the markup term with affiliation as m_d and the markup term from the static model as m_s, we compare these two terms at the solution to the static model:

$$m_d = \frac{-s^0 - \lambda r^{ss}\left(s^j - s^0\right)}{\frac{ds^0}{dp} + \lambda \frac{dr^{ss}}{dp}\left(s^j - s^0\right) + \lambda r^{ss}\left(\frac{ds^j}{dp} - \frac{ds^0}{dp}\right)} \gtreqless -\frac{s^0}{ds^0/dp} = m_s.$$

For a given price, the terms s^0 and ds^0/dp are equivalent across the two models. First, we substitute for r^{ss} in terms of choice probabilities from equation (E.1), as well as its derivative with respect to p. Then, after rearranging terms, we obtain a simple condition relating the levels of the markup terms:

$$m_d \gtreqless m_s \iff -\frac{\partial s^0}{\partial p} \gtreqless -\frac{\partial s^j}{\partial p}. \tag{E.3}$$

A higher value for m_d indicates higher markups and higher prices. Thus, whether markups are higher in the dynamic model depends on the price sensitivity of affiliated consumers relative to shoppers. We can further say that if the denominator of m_d is negative, then $m_d > m_s \iff -\frac{\partial s^0}{\partial p} > -\frac{\partial s^j}{\partial p}$, meaning that markups are higher if affiliated customers are less price sensitive. The sign of the denominator will depend upon the specified demand system.

This is an intuitive result. However, there is a nuanced point to this analysis, stemming from the fact that there is not a direct mapping between our assumption of positive dependence and the condition in (E.3). Given our extension of the logit formulation, $\frac{\partial s^0}{\partial p} = \frac{\partial \delta}{\partial p} s^0(1 - s^0)$ and $\frac{\partial s^j}{\partial p} = (\frac{\partial \delta}{\partial p} + \frac{\partial \sigma}{\partial p}) s^j(1 - s^j)$. Thus, whether or not markups are higher also depends on the derivative of the type-specific shock with respect to price and the relative distance of s_0 and s_j from 0.5 (at which point $s(1-s)$ is maximized). Therefore, steady-state markups may be higher or lower with the presence of consumer affiliation. If we make the additional assumption that affiliated consumer utility is less sensitive to price, i.e., $-\frac{\partial \delta}{\partial p} > -(\frac{\partial \delta}{\partial p} + \frac{\partial \sigma}{\partial p})$, we might expect that markups are higher in the presence of consumer affiliation. However, the results show that it is still ambiguous whether markups are higher in the steady state, as s_j may be close enough to 0.5 relative to s_0 to flip the inequality.

Thus, the presence of positively affiliated consumers may, counter-intuitively, lower the steady-state price, relative to the static model. The intuition for this result is akin to those summarized in Farrell and Klemperer (2007); with dynamic demand and affiliation, firms face a trade-off between pricing aggressively today and "harvesting" affiliated consumers in future periods. In the steady state, our model shows that either effect may dominate.