


Working Paper 19-111

Consumer Inertia and Market Power

Alexander MacKay
Marc Remer



**Harvard
Business
School**

Consumer Inertia and Market Power

Alexander MacKay
Harvard Business School

Marc Remer
Swarthmore College

Working Paper 19-111

Copyright © 2019, 2021, 2022 by Alexander MacKay and Marc Remer.

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Funding for this research was provided in part by Harvard Business School.

Consumer Inertia and Market Power*

Alexander MacKay
Harvard University[†]

Marc Remer
Swarthmore College[‡]

January 29, 2022

Abstract

We study the pricing decision of firms in the presence of consumer inertia. Inertia, which can arise from habit formation, brand loyalty, and switching costs, generates dynamic pricing incentives. These incentives mediate the impact of competition on market power in oligopoly settings. For example, dynamic incentives can limit the equilibrium price effects of a horizontal merger. However, the way that the merger is implemented—whether the merged firm maintains separate brands or consolidates them into a single entity—can have large effects on equilibrium prices in the presence of inertia. We develop an empirical oligopoly model to estimate consumer inertia and dynamic pricing incentives using market-level data. We apply the model to a hypothetical merger of retail gasoline companies, and we find that a static model predicts greater price increases than those obtained while accounting for dynamics.

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

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

*An earlier version of this paper was entitled “An Empirical Model of Consumer Affiliation and Dynamic Price Competition.” 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, US Department of Justice, NYU Stern, and Harvard University for their feedback and suggestions.

[†]Harvard University, Harvard Business School. Email: amackay@hbs.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 reflects both exogenous preferences and state-dependent utility that is affected by past behavior (Heckman, 1981). Consumer state dependence, or inertia, may arise from habit formation, brand loyalty, switching costs, or search. There is a rich empirical literature that establishes the presence of such inertia in a variety of markets.¹ In response to consumer inertia, profit-maximizing firms will internalize the effect of their current price on demand in future periods.

These incentives have important implications for competition in oligopoly settings. Inertia can lead to higher equilibrium prices, as consumers become locked into a particular product or brand over time. We denote the extent to which a firm sets higher prices due to dynamic consumer behavior as *dynamic market power*.² We show that dynamic market power can influence *horizontal market power*—the extent to which competitors constrain prices. Therefore, it is necessary to account for both the dynamic and horizontal dimensions of market power when assessing market competition, evaluating the potential effects of mergers, or designing competition policy.

We study the interaction of these forces using a dynamic oligopoly model where firms compete in prices. We model consumer inertia as a product-specific preference, or “affiliation,” for the most recently purchased product. This formulation nests typical implementations of dynamic consumer behavior, including switching costs and brand loyalty. We use the model to evaluate the impact of consumer inertia on steady-state prices with different degrees of market competition. We then analyze the impact of changes in competition arising from horizontal mergers, while holding fixed the level of inertia. We demonstrate that accounting for consumer inertia is critical to accurately predict the price effects of mergers. Furthermore, as we describe below, whether or not a merged firm consolidates or maintains separate brands, which may not matter in static models, can have substantive implications in the presence of inertia.

Our findings highlight the importance of accounting for dynamic pricing incentives. We develop an empirical model to evaluate these incentives and provide an approach to simulate counterfactuals with dynamic oligopoly pricing. We use counterfactuals to quantify the price effects of inertia and competition in an empirical application. Thus, we create a framework to account for both the dynamic and horizontal dimensions of market power for key policy questions.

Typically, firms and competition authorities employ static empirical models to estimate consumer demand and simulate counterfactual prices. Competition authorities will challenge a merger if the merging firms are expected to increase prices by a significant amount. If con-

¹See the related literature section for examples.

²As consumer inertia can also lead to lower equilibrium prices (Dubé et al., 2009), it is possible that dynamic market power is lower with greater inertia.

sumer inertia is present in the market, these static models omit the first-order incentive of a firm to reduce price to increase future demand. Instead, the static model attributes this effect to less differentiation among competing firms. This can cause static models to substantially over-predict the price effects of a merger. Thus, failing to account for consumer inertia may misstate the potential for horizontal market power and affect merger enforcement.

Competition authorities often analyze mergers in markets that are likely to be characterized by consumer affiliation. For example, in its lawsuit against Swedish Match-National Tobacco, the Federal Trade Commission (FTC) cited strong brand loyalty as a barrier to entry.³ Similarly, the US Department of Justice cited customer switching as an important factor in its case against the UPM-MACtac merger.⁴ In defense of its acquisition of TaxACT, H&R Block cited the importance of dynamic incentives in exerting downward pricing pressure post-merger (Remer and Warren-Boulton, 2014). Both the FTC and DOJ routinely investigate mergers in consumer product markets, where inertia in brand choice has been documented. Consumer inertia is a central feature of many technology markets, such as those served by Amazon, Google, Apple, and Facebook, which are currently under heavy scrutiny by antitrust agencies. Yet, perhaps due to computational complexity or compressed investigative timelines, the supply-side impacts of inertia are seldom quantified in these investigations. We develop a relatively parsimonious model that substantially reduces the computational and data burden of estimating the impact of the dynamic incentives facing firms.

We begin by evaluating theoretical properties of our model using numerical simulations. We demonstrate that affiliation can have a large effect on steady-state prices. In equilibrium, affiliation raises oligopoly prices and often has a greater impact on price than a reduction in competition. In evaluating horizontal market power, we show that accounting for consumer inertia and also the *type* of merger is critical. A merged firm can decide to maintain separate brands or consolidate merged brands into a single entity. In a static setting, this distinction may not be relevant; we show that mergers of either type deliver the exact same outcome for symmetric firms with logit demand. However, these two types of mergers can deliver directionally opposite outcomes in the presence of consumer inertia. These results are driven by supply-side incentives, not biased demand elasticities. These findings highlight the need for empirical models to appropriately represent firms' dynamic pricing incentives, and for counterfactual exercises to precisely define the structure of the but-for world.

We then develop an empirical model of consumer inertia that can be estimated using aggregate, market-level data. The demand model we introduce is a straightforward extension of the standard discrete-choice logit model with myopic consumers. In contrast to the typical random coefficients model, we allow the distribution of unobserved heterogeneity to be state dependent, affected by past purchase behavior. We restrict the random coefficients to affect a

³See, "Commentary on the Horizontal Merger Guidelines," 2006, <https://www.justice.gov/atr/file/801216/download>

⁴Ibid.

single product for each consumer type, corresponding to our notion of product-level affiliation. To disentangle heterogeneity in preferences from state dependence arising from previous purchases, we impose a demand system and rely on the panel structure of our data. Intuitively, after adjusting for cross-sectional fixed effects, consumer inertia is captured by the residual correlation in shares over time (and their relation to prices).

A key advantage of the model is that it can be estimated using aggregate market-level shares and prices, which is the typical data used in demand estimation and merger simulation. Our model allows us to separately identify the choice probabilities for consumers of different unobserved types and states. Thus, we allow for endogenous unobserved heterogeneity through the presence of a serially correlated state variable for each firm. To identify random coefficients that evolve over time, we limit the degree to which static unobserved heterogeneity enters the model. Thus, we give up more flexible static substitution patterns in order to capture substitution patterns that reflect consumer dynamics, while still relying on aggregate data. In our setting, consumer dynamics are impactful. Importantly, the model can be estimated independently of supply-side assumptions, and we can therefore use the estimated demand model to test for forward-looking behavior by firms.

To highlight the importance of accounting for dynamics when predicting firm behavior, we impose a supply-side model of price competition. In contrast to the literature, we invoke relatively weak assumptions about supply-side behavior in order to conduct counterfactual analysis. From the estimated demand model, we obtain the derivative of static profits, which we use to infer the dynamic component of the firms' first-order conditions. We project these estimates onto state variables to construct a reduced-form approximation to the dynamic pricing incentives. This approximation is consistent with a model of Markov perfect equilibrium where firms use limited state variables to forecast their continuation value. We use forward simulations to verify that realized equilibrium payoffs are consistent with these forecasts. We then use the model to evaluate horizontal and dynamic market power in our empirical setting.

As a case study, we apply the model using data from retail gasoline markets. This setting has a relatively homogeneous product, minimal transaction costs, and is not typically thought of as a product to which consumers are particularly loyal. Thus, to the extent that dynamic incentives matter in this market, we might suspect that such incentives are important across many consumer product markets. There is some existing evidence of pricing patterns in retail gasoline that are consistent with consumer affiliation, such as slow-to-adjust cost pass-through (e.g., Lewis and Noel, 2011). Furthermore, retail gasoline has a direct link to antitrust concerns. In June and December of 2017, the FTC challenged Alimentation Couche-Tard's acquisitions of Empire Petroleum Partners⁵ and Holiday,⁶ respectively, on the basis of overlapping retail

⁵<https://www.ftc.gov/news-events/press-releases/2017/06/ftc-requires-retail-fuel-station-convenience-store-operator>

⁶<https://www.antitrustalert.com/2017/12/articles/ftc-developments/the-latest-ftc-challenges-retail-fuel-station-and-convenience-store-transaction-requires-ten-localized-divestitures-in-wisconsin-and-minnesota/>

gasoline stations in a number of states.

We estimate the model using a rich panel dataset of prices, shares, and costs for retail gasoline stations. In this context, the model is best interpreted as one of habit formation or consumer inattention, wherein some consumers return to the gas station from which they previously purchased without considering alternative sellers. We find evidence of strong demand dynamics. We estimate that 62 percent of consumers have the tendency to become affiliated to the brand from which they previously purchased.⁷ The remaining 38 percent are “shoppers” that are unaffected by consumer inertia. Consumers that become affiliated to a brand are not very price sensitive, with an average elasticity of only -0.55 . Shoppers are extremely price sensitive. Though shoppers are a minority of all consumers that purchase gasoline, they are important for disciplining prices in equilibrium. Across all consumer types, firms face an average elasticity of -3.4 .

To evaluate horizontal market power, we perform a merger analysis between two major gasoline retailers and re-compute the price-setting equilibrium in each period. With the dynamic model, we estimate that brand consolidation would increase prices for the merging firms by 2.0 percent post-merger. A static model, on the other hand, predicts an average price increase of 5.4 percent, which would likely result in greater antitrust scrutiny. Therefore, the dynamic incentive to invest in future demand mitigates the increase in horizontal market power obtained after a merger. With the dynamic model, a merger with joint pricing for separate brands results in a price increase of 3.3 percent for the merging firms.

To evaluate dynamic market power, we calculate equilibrium prices if we increase or decrease the share of consumers affected by inertia. An increase of 7 percentage points, which shifts the population affected by inertia from 62 percent to 69 percent, results in prices that are 2.0 percent higher.⁸ The effect on prices and firm profits are similar to the merger effects with brand consolidation. We conclude that a relatively small share of shoppers can reduce prices by as much as the introduction of a major competitor. Thus, we use the model to assess the relative magnitudes of dynamic and horizontal market power.

Prior to estimating the model, we present reduced-form evidence of dynamic demand and dynamic pricing in retail gasoline markets.⁹ Consistent with investment in affiliated consumers, we find that new entrants initially price lower than established firms but raise prices over time. We then examine cost pass-through. Using the data to separate out expected and unexpected costs, we show that firms respond differentially to these two measures. Firms begin raising prices in anticipation of higher costs approximately 28 days prior to an expected cost shock, further indicating the importance of supply-side dynamics in this setting.

⁷Interestingly, 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>

⁸Conversely, a decrease by 7 percentage points results in prices that are 1.7 percent lower.

⁹By dynamic pricing we mean that there are intertemporal spillovers. This should not be confused with static pricing in response to changing market conditions, which is often colloquially referred to as “dynamic pricing.”

The paper proceeds as follows: In Section 2, we introduce the model and present results from numerical simulations. Section 3 presents the data for our empirical application and reduced-form evidence of dynamics. Section 4 provides implementation details and results for dynamic demand estimation. In Section 5, we introduce our approach to evaluate dynamic pricing incentives and construct counterfactuals. Section 6 concludes.

Related Literature

We consider the implications of consumer state dependence on the pricing behavior of firms, building on an empirical literature that includes Dubé et al. (2009). We add to the literature by examining the effect of competition and changes in market structure in such settings. As we show, state dependence can have a large effect on the interpretation of outcomes when studying inter-firm competition. Our analysis of mergers complements theoretical work on dynamic price competition when consumers are habit-forming or have switching costs (see, e.g., Farrell and Shapiro, 1988; Beggs and Klemperer, 1992; Bergemann and Välimäki, 2006). Such features link directly to our notion of consumer affiliation.¹⁰ Our empirical model can be used to assess the impacts of competition using real-world data.

We contribute to the empirical literature that estimates state dependence in consumer preferences. Meaningful switching costs, due to brand loyalty or consumer inertia, have been found in consumer packaged goods (Shum, 2004; Dubé et al., 2010), health insurance (Handel, 2013), and auto insurance (Honka, 2014). Hortaçsu et al. (2017) find that consumer inattention and brand loyalty lead to substantial inertia in retail electricity markets. Bronnenberg et al. (2012) and Eizenberg and Salvo (2015) evaluate the presence of habit formation in soft drink markets. When explicitly modeled, the above papers assume that consumers are myopic when choosing among products, which we also maintain. Conceptually, our demand model shares similar features to that of Dubé et al. (2009) and Dubé et al. (2010). Relative to that model, we give up some flexibility in order to estimate demand with a panel of market-level price and quantity data, rather than requiring individual customer data. These less stringent data requirements, as well as our supply-side innovations, allow our model to be applied in a wide variety of settings where it is typically difficult to account for consumer inertia.

The existing literature has primarily relied on consumer-specific purchase histories to document state dependence, whereas our method allows for the recovery of such state dependence using aggregate, market-level data. For an analysis of inter-firm competition, such data tends to be more readily available. One paper that has used aggregate data to estimate switching costs is Shcherbakov (2016), who provides an intuitive argument for the identification of switching costs from aggregate data in the context of television services.¹¹ One innovation of our frame-

¹⁰There are alternative strategic reasons for dynamic pricing, including experience goods (Bergemann and Välimäki, 1996), network effects (Cabral, 2011), learning-by-doing (Besanko et al., 2018), and search (Stahl, 1989).

¹¹We are aware of two other papers that estimate switching costs using aggregate data: Nosal (2011) and Yeo

work is that we allow for a proportion of consumers that are unaffected by consumer inertia, and we specifically estimate the fraction of consumers of this type.

We contribute to a growing body of empirical models of dynamic demand. Existing work focuses on different contexts that drive dynamic behavior. Hendel and Nevo (2013) consider a model with storable goods and consumer stockpiling. Gowrisankaran and Rysman (2012) and Lee (2013) consider the purchase of durable goods with forward-looking behavior by consumers. In contrast to these papers, we focus on settings with positive dependence in purchasing behavior over time. The literature highlights the issue, common to our setting, that misspecified static models can produce biased elasticities. Hendel and Nevo (2013) point out that this will matter in a merger analysis. We complement this point by providing a case in which the dynamic incentives, rather than biased elasticities, are a primary concern in model misspecification.

For our empirical application, we propose a reduced-form method to approximate the dynamic incentives in supply-side pricing behavior, which allows us to side-step some of the challenges present in the estimation of dynamic games. Compared to value-function approximation methods proposed by Bajari et al. (2007) and Pakes et al. (2007), we rely more heavily on the structure of the demand model and place weaker assumptions on supply-side behavior. We estimate the derivative of the value function directly, circumventing some of the computational challenges of estimating the value function (see, e.g., Farias et al., 2012; Sweeting, 2013) and eliminating a recursive step. We motivate our estimation of this function as either a limited-information equilibrium concept or an approximation to full-information behavior by firms, as in Weintraub et al. (2008) and subsequent work. Our focus on the pricing behavior of firms precludes the use of several developments in the conditional choice probabilities literature, which relies on discrete actions (e.g., Aguirregabiria and Mira, 2007; Arcidiacono and Miller, 2011).

2 A Model of Oligopolistic Competition with Consumer Inertia

We develop a dynamic model of oligopolistic competition with product differentiation where consumers may become affiliated with the firm from which they purchased previously. Affiliation may be interpreted as habit formation, brand loyalty, switching costs, or search. Consumers in the model are myopic in that they maximize current period utility rather than a discounted flow of future utility. This assumption is likely a good fit for retail gasoline markets, where consumers do not choose a gas station anticipating that it will limit their future choice set. Despite this, some consumers are likely to return to the same gas station due to habit formation, brand loyalty, or inattention. Similarly, myopia is also likely a reasonable assumption for other consumer products, including many products purchased at grocery stores, where product

and Miller (2018). These papers have less formal identification arguments than Shcherbakov (2016).

affiliation has been documented.¹²

As detailed below, we introduce consumer dynamics by allowing for endogenous unobserved heterogeneity in a differentiated product demand model. We then place the demand model into a dynamic oligopoly setting. Even though consumers are myopic, key dynamics arise when firms internalize the effect of sales today on future profits through the accumulation of affiliated consumers.¹³ We use the model to numerically and empirically examine the impact of consumer inertia on market power, in general, and in the context of horizontal mergers.

2.1 Demand

We extend the standard logit discrete choice model to allow for unobserved heterogeneity that depends on past purchases. The first assumption below presents a random coefficients utility formulation with myopic choice. The second assumption restricts the random coefficients so that the type-specific utility shock affects only a single product, corresponding to our notation of consumer affiliation. The third assumption places restrictions on the evolution of consumer types over time.

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). When unobserved heterogeneity is exogenous, consumers may be indexed by discrete types.¹⁴ We allow the distribution of preferences to change endogenously over time, therefore, heterogeneity in preferences is represented here by *states*. In any period, each consumer is indexed by a discrete state $z \in Z$.

A consumer i in state z receives the following utility for choosing product j :

$$u_{ijt}(z) = \delta_{jt} + \sigma_{jt}(z) + \epsilon_{ijt}. \quad (1)$$

Consumers receive an additively-separable common component δ_{jt} , a state-dependent shock $\sigma_{jt}(z)$, and an idiosyncratic shock, ϵ_{ijt} . The common component will typically be a function of firm j 's price, and takes the form $\delta_{jt} = \xi_j + \alpha p_{jt}$ in the standard logit model (with $\alpha < 0$). The state-dependent shock, $\sigma_{jt}(z)$, may be also be a function of firm j 's price, if consumers are less sensitive to the price of the product to which they are affiliated.

We denote the probability that a consumer in state z chooses product j as $s_{jt}(z)$. We normalize the utility of the outside good to be zero. Additionally, we define the state 0 consumer to be “unaffiliated” with no state-dependent preference, i.e., $\sigma_{jt}(0) = 0 \forall j$. Given the stan-

¹²For example, see Dubé et al. (2010).

¹³Slade (1998) estimates a model of habit-forming consumers and sticky prices. That model, however, explicitly imposes a cost of price adjustment. Our model does not rely upon a menu cost to explain dynamic price adjustments.

¹⁴The discrete type assumption for the random coefficient model is made elsewhere in the literature, though these types are assumed to be exogenous. See, for example, Berry et al. (2006) and Berry and Jia (2010).

standard assumption of a type 1 extreme value distribution on the utility shock, ϵ_{ijt} , the choice probabilities of consumers are:

$$s_{jt}(0) = \frac{\exp(\delta_{jt})}{1 + \sum_k \exp(\delta_{kt})} \quad (2)$$

$$s_{jt}(z) = \frac{\exp(\delta_{jt} + \sigma_{jt}(z))}{1 + \sum_z \exp(\delta_{kt} + \sigma_{jt}(z))}. \quad (3)$$

The observed share for firm j is given by the weighted average of choice probabilities for consumers in each state: $S_{jt} = \sum_{z=0}^Z r_{zt} s_{jt}(z)$, where r_{zt} is the fraction of consumers in state z .

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 allow for serial correlation in unobservable utility shocks over time.

Assumption 2: Single-Product Affiliation We now place restrictions on the state-dependent demand shocks, $\sigma_{jt}(z)$. We assume that each consumer state corresponds to an affiliation (utility shock) to a single product. Further, we assume that there is a single state corresponding to each product. Thus, a consumer in state j is affiliated to product j . The state-specific demand shocks for other products are zero, i.e. $\sigma_{jt}(z) = 0 \forall z \neq j$.

Thus, we define *affiliation* to be a product-specific state dependence in preferences. We say that a consumer in state j is affiliated with product j , as this consumer has a perceived benefit of $\sigma_{jt}(j)$ relative to unaffiliated (state 0) consumers. The affiliation shock $\sigma_{jt}(j)$ has different interpretations depending on the underlying mechanism:

- *Brand loyalty*: The model has a direct interpretation of brand loyalty when $\sigma_{jt}(j)$ is a positive level shock that reflects an internal benefit for purchasing from the same brand.
- *Switching costs*: The model may be interpreted as a switching cost model when $\sigma_{jt}(j)$ is a level shock representing the costs (physical and psychic) of switching to another brand. This interpretation is empirically indistinguishable from the brand loyalty model because only the relative utilities affect choices in the logit formulation of the discrete choice model.
- *Habit formation*: In the habit formation interpretation, a consumer gets either an extra benefit for repeating earlier behavior or bears a cost for adjusting behavior. $\sigma_{jt}(j)$ represents the net benefit. In contrast to the switching cost model, other aspects of preferences may change. For example, consumers may become less price sensitive to the affiliated product, in addition to realizing a level shock.
- *Search/Inattention*: In the special case where $\sigma_{jt}(j)$ renders affiliated consumers inelastic, the model has a search or inattention interpretation. In this case, the unaffiliated

consumers are those that engage in search and realize full information about the choice set. Affiliated consumers are inattentive and simply buy the previous product. This extends standard search models (e.g., Varian, 1980; Stahl, 1989) by allowing for a mixture of consumers that search and those that do not. The fraction that do not search is endogenous and depends on past prices.

Distinguishing among these different mechanisms lies outside of the scope of this paper but may be important, especially when examining questions about welfare. The brand loyalty model and the switching cost model can have identical outcomes but divergent welfare predictions, as $\sigma_{jt}(j)$ is a net benefit in the former and a net cost in the latter.

Assumption 3: Evolution of Consumer States Our framework allows us to consider both endogenous and exogenous unobserved heterogeneity. We consider unobserved exogenous heterogeneity in a simple form by assuming that there are two latent consumer types. A fraction λ of consumers are affected by state dependence. Each of these consumers becomes affiliated with the product they purchased in the previous period. The remaining $1 - \lambda$ fraction of consumers are unaffected by state dependence. These represent “shoppers” that have no costs of switching among different products over time. This fraction of consumers remain in state 0 regardless of their choices.

Let $\lambda_i = 1$ if consumer i is affected by state dependence and $\lambda_i = 0$ otherwise, and let a_i denote consumer i 's choice in the previous period. We can write the indirect utility for consumer i of choosing product j as:

$$u_{ijt}(z) = \delta_{jt} + \mathbb{1}[j = a_i]\lambda_i\sigma_{jt} + \epsilon_{ijt}. \quad (4)$$

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. The coefficient λ_i captures exogenous unobserved type heterogeneity in types, and the coefficient $\mathbb{1}[j = a_i]$ captures unobserved heterogeneity that is endogenous and depends on past prices.

Based on these assumptions, we can represent the state of each market in each period by the vector $\{r_{jt}\}$, where r_{jt} denotes the share of λ -type consumers affiliated to product j in period t . Since λ -type consumers may be unaffiliated, $\sum_{j=0}^J r_{jt} = 1$. λ -type consumers that choose the outside good transition into state 0 for the next period but maintain their latent type. Thus, the share of consumers that are shoppers depends on previous period choices and the underlying distribution of types.

2.2 Supply

We assume that firms set prices to maximize the net present value of profits. We restrict attention to Markov perfect equilibria. Current-period profits are a function of shares. Firm j 's

aggregate share across consumer types and states is:

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

Thus, a firm's total share of sales can be written as a weighted sum of its share of unaffiliated consumers, $s_{jt}(0)$, and affiliated consumers, $s_{jt}(z)$. Note that firm j will make sales to consumers affiliated to other firms $z \neq j$, but the probability that such consumers will choose firm j is strictly lower than the choice probability of an unaffiliated consumer when the utility shocks $\{\sigma_{jt}(z)\}$ are positive.

Assumption 4: Competition in Prices We assume that 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 in each period is summarized by marginal costs, c_t , the distribution of affiliation across consumers, r_t , and other variables that are captured by the vector, x_t , such as expectations about future costs. Entry is exogenous. The firm's objective function can be summarized by the Bellman equation:

$$V_j(c_t, r_t, x_t) = \max_{p_{jt}|p_{-jt}} \{(p_{jt} - c_{jt})S_{jt} + \beta E(V_j(c_{t+1}, r_{t+1}, x_{t+1})|p_t, c_t, r_t, x_t)\}. \quad (6)$$

Prices in each period optimize the sum of current-period profits $(p_{jt} - c_{jt})S_{jt}$ and the continuation value. Both of these components depend upon marginal costs and the distribution of consumer states, r_t . Thus, when the perceived continuation value is non-zero, firms anticipate how price affects the future distribution of consumer states and also the impact of future changes to marginal costs. Note that the state space does not include previous period prices. We therefore exclude strategies that depend directly upon competitors' historical prices, such as many forms of collusion.

Assumption 5: Expectations Consistent with the Markov perfect framework, we make the relatively weak assumption 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 is to directly estimate a reduced-form model of (the derivative of) the continuation value. We describe this approach in more detail in Section 5.

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

2.3 Theoretical and Numerical Analysis

To develop an understanding of pricing incentives in markets with consumer inertia, we consider a deterministic setting where marginal costs are constant. We use numerical methods to analyze steady-state prices in an oligopoly game with Bertrand price-setting behavior.¹⁵ For the numerical exercise, we specify the utility parameters from equation (1) as $\delta_{jt} = \xi_j + \alpha p_{jt}$ and $\sigma_{jt}(j) = \bar{\xi}$, allowing for the affiliation shock to affect only the utility level. Thus, the utility of a consumer i in state z for product j is as follows:

$$u_{ijt}(z) = \delta_{jt} + \mathbb{1}[j = z] \lambda_i \bar{\xi} + \epsilon_{ijt}. \quad (7)$$

Here, $\mathbb{1}[j = z]$ takes a value of one if a consumer that purchases product z in period $t - 1$ purchases product $j = z$ in period t , and zero otherwise. Consumers subject to affiliation always have a value of $\lambda_i = 1$, and λ_i is zero otherwise. Assuming the error term, ϵ_{ijt} , follows the type-1 extreme value distribution, we get the market share specified in equations (2) and (3) in the previous section, which reduces to the standard logit model if $\bar{\xi} = 0$.

Each firm k sells a set of products, $j \in J_k$, and maximizes the expected discounted value of profits. Therefore, firm k 's value function takes the following form:

$$V_k(r) = \max_{p_k | p_{-k}} \pi_k(p, r) + \beta V_k(r'). \quad (8)$$

Here, p and r are vectors of prices and affiliated customers, respectively, and r' is a vector specifying each firm's affiliated customers in the next period. In accordance with the model above, an element of r' is $r'_k = f(p, r) = \sum_{z \in 0, J} r_{zt} s_{jt}(z) = \frac{1}{\lambda} (S_{jt} - (1 - \lambda) s_{jt}(0))$. Static profits are $\pi_k(p, r) = \sum_{j \in J_k} (p_j - c_j) \cdot s_j(p, r)$. We drop the expectations operator, as the only source of uncertainty in the model is the realizations of marginal costs, which are fixed in the steady state. To find the steady-state prices and affiliated shares for each firm, we focus on Markov perfect equilibrium.¹⁶ Firm k 's profit-maximizing first-order condition is then:

$$\sum_{l \in J_k} \frac{\partial \pi_l}{\partial p_j} + \beta \frac{dV_k(r')}{dr'} \frac{dr'}{dp_j} \quad \forall j \in J_k = 0. \quad (9)$$

Next, we specify the derivatives of equation (8) with respect to r and evaluate them at the prices that solve each firm's first-order condition, which will be the prevailing prices at the steady state. These derivatives, in conjunction with the steady-state condition, $\frac{dV'}{dr'} = \frac{dV}{dr}$, yield the following system of equations:

¹⁵We provide a theoretical analysis of the monopolist's problem in the Appendix.

¹⁶Although we do not prove that the equilibrium is unique, the simulation results support there being a single steady-state equilibrium.

$$\underbrace{\frac{dV_k(r)}{dr}}_{J \times 1} = \underbrace{\left[\frac{\partial \pi_k}{\partial p} \frac{dp}{dr} + \frac{\partial \pi_k}{\partial r} \right]}_{J \times 1} \underbrace{\left[I - \beta f_p(p, r) \frac{dp}{dr} - \beta f_r(p, r) \right]^{-1}}_{J \times J}. \quad (10)$$

In this equation, $\frac{\partial \pi_k}{\partial p}$, $\frac{\partial \pi_k}{\partial r}$, $f_p(p, r)$, $f_r(p, r)$ are known, conditional on values of p and r . The remaining unknowns are $\frac{dp}{dr}$ and $\frac{dV_k(r)}{dr}$. To solve the model, we impose the steady-state condition governing the evolution of affiliated customers, $r' = r$. The full set of steady state conditions, provided by equation (10) and $r' = r$, allow us to solve for steady-state prices and shares, conditional on the $J \times J$ derivative matrix, $\frac{dp}{dr}$. The values of $\frac{dp}{dr}$ are determined by the model. In our simulations, we solve for these values numerically using a local approximation method. For additional details, see the Appendix.

For the oligopoly analysis, we simulate symmetric markets with three single-product firms. In this setting, symmetry is imposed by restricting the utility parameters, ξ_j , $\bar{\xi}$, and marginal cost to be the same across all products. Consistent with Dubé et al. (2009), equilibrium prices may be increasing or decreasing for different levels of affiliation. With greater affiliation, firms face less elastic demand, but the incentive to invest in future demand tends to increase. Whether or not prices increase with affiliation depends on the relative weights on these forces and underlying market parameters. For examples, see Figure 6 in the Appendix, which plots equilibrium prices for monopolists and three firms for different draws of the utility parameters. Importantly for our context, affiliation may have differential effects depending on market structure, so that the price effect of a merger may be larger or smaller with greater consumer dynamics.

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 the pre-merger equilibrium prices is given by $\frac{\partial \pi_b}{\partial p_j} = \sum_{l \in \{J_k, b\}} \frac{\partial \pi_l}{\partial p_j} - \sum_{l \in J_k} \frac{\partial \pi_l}{\partial p_j}$. In the case of substitutes, $\frac{\partial \pi_k}{\partial p_j}$ is greater than zero and the firm raises prices.

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

$$\text{Change in dynamic first-order condition: } \frac{\partial \pi_b}{\partial p_j} + \beta \left(\frac{d\tilde{V}_k(r')}{dr'} - \frac{dV_k(r')}{dr'} \right) \frac{dr'}{dp_j}, \quad (11)$$

where \tilde{V}_k now also incorporates the discounted flow of profits from product b . The first term, $\frac{\partial \pi_k}{\partial p_j}$, is equivalent to the change in the static first-order condition, though it is evaluated at different equilibrium prices. When the share of affiliated consumers has a greater marginal effect on future profits $\left(\frac{d\tilde{V}_k(r')}{dr'} > \frac{dV_k(r')}{dr'} \right)$ and the pre-merger investment incentives outweigh the harvest incentives, then the second term is negative and consumer dynamics have a muting effect on potential price increases. The reverse can also occur. The characteristics of the market and also the type of merger determine how consumer dynamics affect post-merger price increases.

2.3.1 Dynamic and Horizontal Market Power

To summarize the relationship between consumer inertia and market power in our model, we use simulations to decompose the potential impacts of dynamic and horizontal market power. To measure dynamic market power, we compare the three-firm oligopoly price with consumer affiliation to a baseline price where consumers have no state dependence, but the markets are otherwise identical. To measure horizontal market power, we compare the three-firm oligopoly price to the price that prevails following a merger between firms one and two.¹⁷

In the presence of consumer inertia, it is important to define precisely the implementation of a horizontal merger. We consider two types of mergers. The first type of merger unites pricing control of two products under a single firm and the merged firm maintains these as separate products. We refer to this as a *joint pricing* merger. This is a common setting in many differentiated product mergers when the combined firms maintain separate brands/products. The second type of merger consolidates two products under a single brand and effectively offers a single product after the merger. We refer to this type of merger as *brand consolidation*. To implement a brand consolidation merger, we need to make an assumption about how much utility consumers of the product removed from the market will receive from buying the remaining brand. We assume that, at pre-merger prices, the consolidated brand will have the same combined share of unaffiliated customers as the separate pre-merger brands.¹⁸ This type of merger typically occurs in settings where retail location is important and the post-merger firm rebrands all locations under one brand. This is often the case in retail gasoline mergers, but also occurs in other industries, such as wireless phone service (T-Mobile/Sprint) and airlines (American Airlines/USAir).

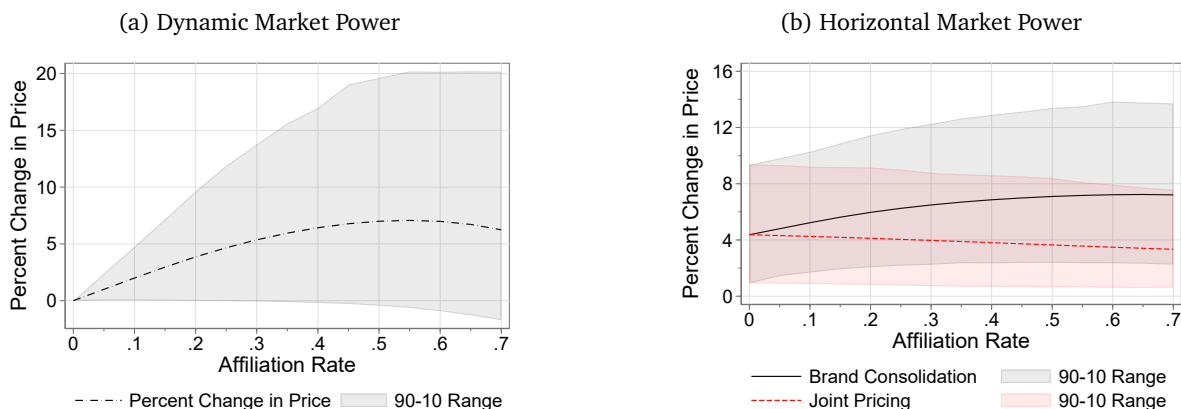
To provide a more generalizable analysis, we attempt to simulate data from the support of parameters that produce reasonable outcomes for margins and shares. We employ a “shotgun” approach, generating simulations with many different parameters and selecting only the markets that meet certain criteria. We first take Halton draws of the demand parameters such that $\xi \in [0, 10]$, $\bar{\xi} \in [0, 10]$, $\alpha \in [-10, 0]$, and set each firm’s marginal cost to one. For each draw of these demand parameters, we construct three-firm markets for $\lambda \in \{0.05, 0.1, 0.15, \dots, 0.70\}$. 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 margins between 0.05 and 0.75.¹⁹ Finally, to avoid composition affects, we only analyze markets with demand parameters that converged for all values of λ . The data generating process yields 6,566 markets whose parameters are

¹⁷Note that we use a (competitive) oligopoly price as a baseline, rather than price equal to marginal cost. Also, as marginal cost is constant across simulations, using price as a measure of market power is equivalent to the commonly used price-cost metric.

¹⁸To implement this assumption, we make an adjustment to the utility derived from consuming the post-merger brand. See the appendix for more details.

¹⁹The range for each parameter is selected such that parameter values toward each edge of the range result in outcomes that fall above or below our share and margin criteria.

Figure 1: Potential Price Effects



Notes: Panel (a) displays the mean percent price increase and for a three-firm oligopoly above 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 $\lambda \in \{0.05, 0.1, 0.15, \dots, 0.70\}$, or 6,566 markets in total.

summarized in Appendix Table 10.

In Figure 1, we plot the effects of affiliation and reduced competition on prices. The plots employ simulation results from the 469 baseline parameter values of $\xi, \bar{\xi}$, and α that converged for all $\lambda \in [0.05, 0.70]$, or 6,566 markets. Panel (a) measures dynamic market power 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. At lower values, a marginal increase in λ raises the “harvest” incentive more than the “invest” incentive, leading to higher prices. However, above some threshold, an increase in λ yields a relatively stronger incentive to invest in future demand, and prices fall. 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 the transparent area. When $\lambda > 0.4$, the 90th percentile market exhibits dynamic market power resulting in prices more than 15 percent higher than a static model.

In panel (b), we plot the price 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 . This is in part due to how we specify brand consolidation mergers, as we assume high customer retention at the consolidated

Table 1: 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 Margin	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	776	76	339	1727	2696	6566
Δ HHI	712	517	51	226	1151	1797	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
Brand Consolidation 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	67.3	99.9	0.1	12.5	74.3	792.0	6566
Prediction Bias: Brand Consolidation	-1.17	1.46	-8.63	-1.77	-0.22	4.34	6566
Prediction Bias (pctg.): Brand Consolidation	-19.7	24.0	-76.0	-35.9	-4.2	181.0	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: Margin 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.

brand. Still, there is large overlap in the 90-10 percentile range across the two merger types, demonstrating that both can enable comparable levels of horizontal market power. Another interesting feature of panel (b) is that horizontal market power slightly decreases, on average, in joint pricing mergers, but it increases in brand consolidation mergers. This highlights that the interaction between customer affiliation and a decrease in competition depends critically on the mechanism through which competition is reduced.

2.3.2 Consumer Inertia and Mergers

To further demonstrate the effect of consumer inertia on price competition, we examine more closely the impact on prices of horizontal mergers. We proceed by again using three symmetric, single-product firms as a baseline, and simulate both joint pricing and brand consolidation mergers. Thus, we examine how dynamic incentives affect steady-state prices and the unilateral price effects arising from mergers. We also analyze the bias that arises when a static model is calibrated to data generated by the dynamic affiliation model and used for merger simulation. We find that ignoring the presence of consumer affiliation tends to overestimate the price effects of a merger. The bias is much more pronounced, however, in joint pricing mergers.

For each parameterized market, we solve for the steady-state prices. Then, we simulate a merger between firms 1 and 2, leaving firm 3 as the non-merging firm. Table 1 summarizes the pre-merger equilibrium and the unilateral price increases of merging and non-merging

firms. The “average” market is one that, a priori, would typically raise moderate concern from the US antitrust agencies; 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. We use these mean values to calculate the “Upward Pricing Pressure” index, $\frac{0.17}{1-0.17} \cdot (1.26 - 1) = .053$, which is just over the threshold that may trigger an investigation.²⁰ The average percentage merger price effect is 3.84 percent and 6.49 percent in joint pricing and brand consolidation mergers, respectively, either of which may raise antitrust concerns. 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 the pricing incentives of firms in the consumer affiliation model.

Tables 11 and 12 in the Appendix provide results from regressing the demand parameters on effects of the mergers. On average, increasing the rate of affiliation (λ) and the strength of the affiliation ($\bar{\xi}$) tends to increase pre-merger prices, but their impact on the effect of a merger is dependent upon the type of merger. The proportion of affiliated customers and the strength of affiliation decrease prices in joint pricing mergers, but work to increase prices in brand consolidation mergers. However, these relationships do not hold in every instance, and may interact in interesting ways. Thus, to fully understand the impact of horizontal mergers in markets with consumer inertia, it is important to model the market under consideration.


2.3.3 Model Misspecification in Merger Simulation

The above results suggest that affiliation has implications for counterfactual exercises, such as merger simulation. Failing to account for consumer inertia will result in biased elasticities from incorrect first-order conditions. Antitrust agencies often infer elasticities from markups calculated using accounting data (see Miller et al., 2013), which omit the dynamic incentive of firms. In addition to generating incorrect elasticities, failing to account for the dynamic incentives in first-order conditions can have large direct effects on post-merger predictions.

To more precisely analyze the impact of model misspecification, we measure the implications of failing to account for consumer affiliation when calibrating demand and simulating a merger. To do so, we consider the following hypothetical scenario. The true underlying model is the three-firm market with consumer state dependence due to affiliation. A practitioner observes each firms’ pre-merger prices, marginal costs, and aggregate market shares (rather than separately observing its unaffiliated and affiliated shares). This data is then used to recover the demand parameters of the standard logit model, and then the price effects of a merger are simulated. We perform this experiment for each of the numerically generated markets, and consider both joint pricing and brand consolidation mergers.

²⁰See, 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 2: Distribution of Merger Prediction Bias



Notes: Prediction bias is defined as the 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.


In a symmetric oligopoly, the joint pricing and brand consolidation mergers we consider will have identical price effects when demand is characterized by the static logit model. We provide a proof in Appendix B. Thus, precisely modeling the structure of a horizontal merger may have heightened importance in markets characterized by consumer inertia.

Table 1 summarizes the prediction bias, defined as the static prediction minus the “true” affiliation price effect. The average prediction bias is 1.47 and -1.17 percentage points for joint pricing and brand consolidation mergers, respectively. Scaled by the magnitude of the true price effect, the mean bias is 67.3 and -19.7 percent for the two types of mergers. Thus, for both types of mergers, incorrectly assuming static demand will lead to substantially biased predictions.

Figure 2 plots the distributions of bias across all 6,566 markets. Panel (a) depicts the distribution of bias for joint pricing mergers. In every instance the static model over predicts the true dynamic effect. In a few markets that did not meet our inclusion criteria, the static model under predicted the dynamic effect, so an upward bias will not always occur. However, our results suggest that a static logit demand model incorrectly calibrated to a demand with consumer inertia tends to exhibit upward bias. On the other hand, the bias in brand consolidation mergers may be positive or negative. While static models are more likely to under-predict the true dynamic price effect, almost 25 percent of simulations resulted in over-predictions. These findings again highlight the importance of properly accounting for dynamics when simulating the price effects of mergers. Moreover, we find that the sign of the bias is only systematic in joint pricing mergers.

Figure 3 relates the simulation price effects to the pre-merger market share of each sym-

Figure 3: Price Increase by Market Share



Notes: The lines depict a local polynomial regression of firm 1's merger percent price change on its pre-merger market share.

metric firm. To generate the graph, we run a local polynomial regression of the merger price effect on the average 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) a misspecified static logit model (which generates the same price effect, regardless of the merger type). In line with intuition, the price effect of a merger is increasing with pre-merger market shares. Interestingly, we find that the magnitude of the bias increases with market share with respect to joint pricing mergers, but decreases with respect to brand consolidation mergers. As mergers with high pre-merger market shares tend to be of greater concern to antitrust authorities, this finding highlights that the cost of misspecification may be low if the merger results in brand consolidation, but upward bias is more likely if only pricing control is expected to change.

The “harvesting” and “investment” incentive present in the dynamic model, which is absent in the static model, can explain the biases depicted in Figure 3. With respect to joint pricing mergers, the static model over-predicts the dynamic model regardless of the pre-merger market shares. Thus, as the dynamic incentive to invest in future demand increases, the upward bias of the static model becomes greater in magnitude. Interestingly, the static model is calibrated to be more elastic, on average, than the share-weighted elasticity in the dynamic model (-4.74 vs. -3.86). More elastic demand generates smaller merger price effects in the logit model. Thus, biased predictions from the static model arise primarily from the omission of dynamic incentives to invest in future demand, rather than a biased elasticity or mean utility parameters alone.

In rough terms, margins are determined in the static model by a combination of market elasticities and the degree of competition within the market. The dynamic model has a third factor disciplining margins: the dynamic incentive. When this factor is omitted in the estimation of a static model, greater weight is attributed to the degree of competition (i.e., rival products

are considered to be closer substitutes). This provides some intuition as to why a static model will tend to predict greater unilateral price effects.

Yet, in brand consolidation mergers, the static model tends to under-predict the true price effect. This is possible when brand consolidation mergers present the firm with a stronger “harvesting” incentive relative to joint pricing mergers. The post-merger firm is able to transfer its loyal (i.e., inelastic) customers to a single product. This presents the firm with ability to raise prices to a higher level than when separate products are maintained. In the misspecified static model, the firm is responding to the “average” elasticity of its customers whereas, in the dynamic model, the firm internalizes the different types of consumers into its pricing incentive. Since affiliated customers are much less elastic, this presents a greater opportunity, relative to a static model, to increase price.

These results highlight the benefit of an empirical model that can account for consumer dynamics, which we pursue in the following section. However, for certain applications, an informal analysis of consumer dynamics may provide a useful indication of static model bias. Our results show how price predictions from static merger simulations could be revised downward when only pricing control is expected to change in a merger. In Appendix A.6, we provide more detail on how dynamics affect simulation bias. Therein, we again find that the relationship between the strength of dynamics and bias is a function of the type of merger under evaluation.

3 Reduced-Form Evidence of Dynamics

To motivate the empirical application, we provide evidence of dynamic demand and dynamically adjusting retail gasoline prices. A host of previous studies have found that retail gasoline prices may take multiple weeks to fully incorporate a change in marginal cost. For a review, see Eckert (2013). One innovation of our study is that we use separate measures of unexpected and expected costs to see if, consistent with forward-looking behavior, firms respond differentially to these two types of costs.

3.1 Data

The analysis relies upon daily, regular fuel retail prices for nearly every gas station in the states of Kentucky and Virginia, which totals 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 discounts per-gallon that are likely fixed over the course of a year. The data ranges from September 25th, 2013 through September 30th, 2015. The data was obtained directly 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).

Table 2: Regressions with Share as the Dependent Variable

	(1)	(2)	(3)	(4)	(5)
Price	0.009*** (0.001)	0.000** (0.000)	0.004 (0.002)	-0.005 (0.004)	-0.073*** (0.017)
Lagged Share		0.973*** (0.001)	0.963*** (0.001)	0.554*** (0.002)	0.628*** (0.003)
Price Squared			-0.000 (0.000)	0.001 (0.001)	0.010*** (0.003)
Comp. Price (Mean)			-0.004*** (0.001)	-0.002 (0.001)	-0.108** (0.045)
Comp. Price (SD)			-0.000 (0.001)	0.003* (0.001)	0.086*** (0.024)
Comp. Stations			-0.000*** (0.000)	-0.000*** (0.000)	-0.004*** (0.000)
Num. Stations			0.000*** (0.000)	0.004*** (0.000)	
Num. Brands			-0.001*** (0.000)	-0.003*** (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

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

OPIS also supplied the market share data, which is a standard dataset used by industry participants to track shares.²¹ The data is specified at the weekly, county/gasoline-brand level.²² Due to contractual limitations, OPIS only provided each brand's share of sales, not the actual volume. Thus, 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). To account for differences in demographics across markets, we merge the data with measures of income and population density from the American Community Survey.

3.2 Dynamic Demand: Correlation in Shares Over Time

Though ultimately the importance of demand-side dynamics in the data will be estimated by the model, it is informative to examine the reduced-form relationships between key elements. The dynamic model developed in the previous section 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 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 that are challenging to explain with contemporaneous variables.²³

3.3 Dynamic Pricing

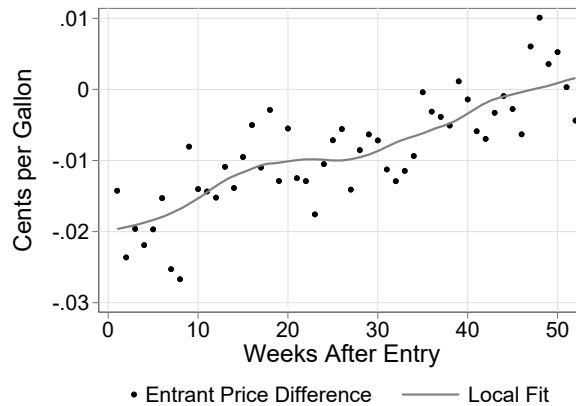
We now present reduced-form evidence of dynamic pricing. Consistent with a model where firms accumulate affiliated consumers over time, we find that new entrants price lower relative to established competitors in the same market, and that this discount dissipates over time. Second, we examine cost pass-through and show that firms are slow to adjust to marginal cost changes. Moreover, firms anticipate expected changes in future costs by raising prices in advance of the change. In the presence of consumer affiliation a firm will change its current price in response to an expected future cost change, as it affects the current value of investing

²¹The data are calculated from "actual purchases that fleet drivers charge to their Wright Express Universal card." Though this may not be a perfect proxy for shares across all driver types, the fact that gasoline brands use it to track the shares of rivals makes it suitable for our purposes.

²²In some instances, the brand of gasoline may differ from the brand of the station. For example, some 7-Eleven stations in the data are identified as selling Exxon branded gasoline.

²³We 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.

Figure 4: 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.

in future demand. The ability to separately estimate the response to expected and unexpected costs is a key innovation of our study.

3.3.1 Dynamic Pricing of New Entrants

When forward-looking firms price to consumers that may become affiliated, there is an incentive to initially offer prices below the static optimum. In this setting, we expect a new entrant, all else equal, to initially price below its competitors. As the new entrant builds up its share of affiliated customers, its prices will gradually converge to its competition.

We test for and find evidence consistent with this dynamic pricing pattern in the data. To perform the analysis, we first identify a set of new entrants, defined as a gas station whose first price observation is at least six weeks after the start of the data and does not exit in the remainder of the sample. To ensure there is sufficient data and to control for composition effects in the analysis, we limit the set of entrants to those with at least one year of post-entry price data. Using this filter, we identify 193 entrants. We compare prices for entrants to the 594 stations in those counties that are present for the entire sample.

Figure 4 depicts the average difference between an entrant's price and all other stations' prices in the same county, sorted by the number of weeks after entry. The figure demonstrates that gas stations enter with a price that is, on average, two cents per gallon less than incumbents' prices. Entrants' prices then slowly converge over time to the market average. For the first 8 weeks following entry, new entrant prices are on average 2.1 cents per gallon lower than incumbents' (standard error: 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, based on our empirical demand estimates, are economically meaningful for attracting unaffiliated

consumers. This pattern is consistent with a profit-maximizing firm building up an affiliated customer base over time, and raising its price to a gradually less elastic set of consumers.

3.3.2 Cost Pass-through

To highlight the temporal component of cost pass-through, we separately estimate how gas stations react to expected versus unexpected cost changes. Beyond motivating the structural model, these results also demonstrate the importance of capturing firms' anticipated price responses when estimating cost pass-through rates. For example, to analyze how much of a tax increase firms will pass-on to consumers, it is imperative to recognize that firms may begin to adjust their prices prior to the tax increase being enacted; failure to account for this response may lead to underestimating pass-through rates.

We construct our measure of expected cost by using gasoline futures and current wholesale costs to project 30-day-ahead costs. Unexpected costs represent deviations from this projection.²⁴ We incorporate the main components of marginal costs for retail gasoline, which include the wholesale cost of gasoline and the per-unit sales tax. We estimate the following model:

$$p_{nt} = \sum_{s=-50}^{50} \beta_s \hat{c}_{n(t-s)} + \sum_{s=-50}^{50} \gamma_s \tilde{c}_{n(t-s)} + \sum_{s=-50}^{50} \phi_s \tau_{n(t-s)} + \psi_n + \varepsilon_{nt}. \quad (12)$$

Here, p_{nt} is the price observed at gas station n at time t . $\hat{c}_{n(t-s)}$ and $\tilde{c}_{n(t-s)}$ are the expected and unexpected wholesale costs observed with lag s , and $\tau_{n(t-s)}$ is the state-level sales tax.²⁵ Using the estimated coefficients on the cost measures, we construct cumulative response functions to track the path of price adjustment to a one time, one unit cost change at time $t = 0$. We incorporate 50 leads and lags to capture the full range of the dynamic response. We focus our results on unexpected and expected costs, as we do not have enough tax changes in our data to estimate a consistent pattern of response.

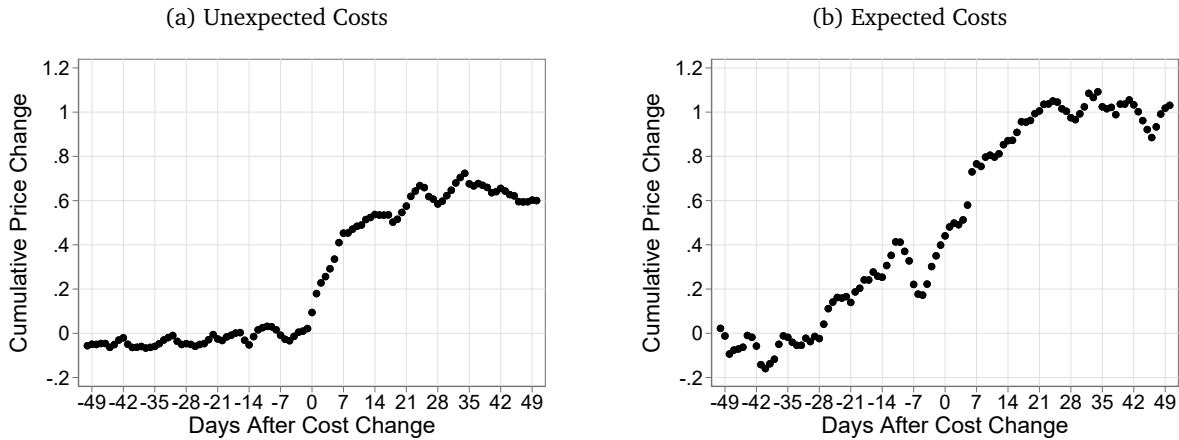
Figure 5 plots the cumulative response functions for unexpected and expected costs. Panel (a) displays the results for unexpected costs. Prices react suddenly and quickly at time zero, but it takes about four weeks for the prices to reach the new long-run equilibrium. Estimated pass-through peaks at 0.72 after 34 days, with an average of 0.64 over days 21 through 50.

Panel (b) displays the cumulative response function for expected costs. Notably, firms begin to react to expected costs approximately 28 days in advance, with a relatively constant adjustment rate until the new long-run equilibrium passthrough is reached 21 days after the shock. The estimated pass-through averages 1.01 over days 21 through 50. Though the total duration of adjustment is longer compared to the unexpected cost shock, the firm incorporates the cost

²⁴For details, see Section C.1 in the Appendix.

²⁵To more easily incorporate future anticipated costs into the regression, we do not present an error-correction model (Engle and Granger, 1987), which is commonly used to estimate pass-through in the retail gasoline literature. As a robustness check, we estimated the price response to expected and unexpected costs using the error-correction model, and we found very similar results.

Figure 5: 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 (12).

more quickly after it is realized. This coincides with substantial anticipation by the firm; the price already captures about 40 percent of the effect of the expected cost shock the day before it arrives.²⁶

Thus, a reduced-form analysis of pricing behavior shows that retail gasoline prices adjust slowly to changes to marginal cost, and also that price changes anticipate expected changes in marginal costs. These patterns are consistent with forward-looking behavior by firms and dynamic demand arising from consumer affiliation. Readers might wonder about the relevance of asymmetric pricing, i.e., whether the price response is the same for positive and negative cost shocks. In robustness checks, we find little evidence of asymmetry. Furthermore, in our data, we do not find evidence of Edgeworth price cycles.

4 Empirical Application: Demand Estimation

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

²⁶A striking result from these estimates is the difference in the long-run pass-through rates. Expected costs experience approximately “full” pass-through – a cost increase leads to a corresponding price increase of equal magnitude. On the other hand, unexpected costs demonstrate incomplete pass-through, moving about only 64 cents for each dollar increase in cost.

these results to consider a merger between large gasoline retailers.

4.1 Identification

We discuss the identification argument in three parts. First, we show that the structure of the model is sufficient to identify the unobserved distribution of choices (i.e., the vector $\{s_{jt}(z)\}$) conditional on observed shares S_{jt} , the share of consumers subject to state dependence λ , and the strength of the affiliation shock $\sigma_{jt}(z)$. Second, the vector $\{s_{jt}(z)\}$ allows us to recover the mean utility for unaffiliated consumers and estimate the static demand parameters. Third, we discuss the assumptions that allow us to identify the dynamic parameters.

Identification of Type-Specific Choices

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 $\{s_{jt}(z)\}$ and depends on the distribution of affiliated consumers for each product $\{r_{jt}\}$. Observed shares are determined by the following:

$$S_{jt} = (1 - \lambda)s_{jt}(0) + \lambda \sum_{z=0}^J r_{zt}s_{jt}(z). \quad (13)$$

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):

$$\ln s_{jt}(0) - \ln s_{0t}(0) = \delta_{jt} \quad (14)$$

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

$$\ln s_{jt}(z) - \ln s_{0t}(z) = \delta_{jt} + \sigma_{jt}(z). \quad (15)$$

To show identification, we combine equations (14) and (15) to obtain the following expres-

sions:

$$s_{jt}(0) = \left(\frac{s_{0t}(0)}{s_{0t}(j)} - 1 \right) \frac{1}{\exp(\sigma_{jt}(j)) - 1} \quad (16)$$

$$s_{jt}(z) = s_{0t}(z) \frac{s_{jt}(0)}{s_{0t}(0)} \cdot \exp(\sigma_{jt}(z)). \quad (17)$$

Thus, we show that 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 $1 - \sum_k s_{jt}(0) - s_{0t}(0) = 0$ and the observed share equations, which provide the other J restrictions:

$$S_{jt} = (1 - \lambda) \left(\frac{s_{0t}(0)}{s_{0t}(j)} - 1 \right) \frac{1}{\exp(\sigma_{jt}(j)) - 1} \quad (18)$$

$$+ \lambda \sum_{z=0}^J r_{zt} s_{0t}(z) \frac{s_{jt}(0)}{s_{0t}(0)} \cdot \exp(\sigma_{jt}(z))$$

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

$$r_{jt} = \sum_{z=0}^J r_{z(t-1)} s_{j(t-1)}(z) \quad (19)$$

$$\implies r_{jt} = \frac{1}{\lambda} (S_{j(t-1)} - (1 - \lambda) s_{j(t-1)}(0)) \quad (20)$$

Given an initial value $\{r_{j0}^*\}$, equation (20) can be used to iteratively identify the future values of the state. We discuss the choice of this initial state vector in our empirical application. Therefore, the unobserved state-dependent choice probabilities $\{s_{jt}(z)\}$ are identified conditional on λ and $\{\sigma_{jt}(z)\}$, i.e., the parameters governing unobserved heterogeneity.

Identification of Static Demand Parameters

We now allow for the observation of multiple markets, which are denoted with the subscript m . From equation (14), we obtain the utility of the unaffiliated (type 0) consumer in each market, $\{\delta_{jmt}\}$. This is analogous to recovering the mean product utility as in (Berry et al., 1995). We make the standard assumption that the utility is linear in characteristics:

$$\delta_{jmt} = \alpha p_{jmt} + \pi (p_{jmt} \times Income_{jmt}) + X_{jmt} \gamma + \eta_{jmt}. \quad (21)$$

The utility depends on price, p , and the interaction of price with market-average income, both of which are endogenous. The exogenous covariates, X , may contain multi-level fixed

effects. With valid instruments for p and $(p \times Income)$, these linear parameters are identified using standard instrumental variables arguments.

Identification of Dynamic Demand Parameters

We have so far shown exact identification of static demand parameters conditional on the parameters governing unobserved heterogeneity. To identify λ_m and $\sigma_{jt}(j)$, we need to employ additional moments. We use fixed effects to model unobserved serial correlation in demand, and we calculate the residual demand innovations, η_{jmt} , after accounting for these fixed effects. Specifically, we include fixed effects that capture aggregate period-specific demand shocks, product-specific persistent demand, and market-specific seasonal patterns. We then assume that the residual demand innovations are uncorrelated over time. We assume that $Corr(\eta_{jmt}, \eta_{jm(t+1)}) = 0$ holds within each market, which provides us with more than sufficient moments (241) to identify our parameters.²⁷ The parameters are pinned down by the patterns of serial correlation in the data and the entry of new brands.

In the context of our model, demand may be serially correlated due to persistent brand-market preferences, aggregate period-specific demand shocks, and market-specific seasonal patterns—all of which are captured by fixed effects. We attribute the residual serial correlation to consumer inertia, where future demand is shaped by the pricing decisions of firms. Thus, we load the systematic autocorrelation in residual demand innovations to the endogenous response of consumers, rather than treating such correlation as a feature of an exogenous stochastic process. Our assumption would be violated if brands systematically realized brand-specific transitory demand shocks within local markets. Depending on the degree to which this is the case, our results may be thought of an “upper-bound” on the impact of consumer inertia.

In our application, we parameterize the dynamic parameters λ_m and $\sigma_{jmt}(j)$ as follows:

$$\lambda_m = \frac{\exp(\theta_1 + \theta_2 Income_m + \theta_3 Density_m)}{1 + \exp(\theta_1 + \theta_2 Income_m + \theta_3 Density_m)} \quad (22)$$

$$\sigma_{jmt}(j) = \bar{\xi}. \quad (23)$$

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. The share of shoppers may be lower in areas with more affluent consumers and more congestion. This specification allows for the possibility that consumer characteristics, as captured by income and population density, affects the prevalence of consumer inertia. We assume that affiliated customers receive a constant level shock to utility $\bar{\xi}$.

²⁷One could construct related moments by using lagged prices as instruments, under the assumption that the prices are uncorrelated with the innovation in the demand residual.

Separate identification of λ_m and $\bar{\xi}$ is made possible by the structure of the model. λ_m , the share of consumers that become affiliated, does not depend on price, whereas the impact of $\bar{\xi}$ on shares does. As can be seen by examining equations (14) and (15), a change in price affects δ_{jmt} , which shifts the relative choice patterns, holding fixed $\bar{\xi}$. Intuitively, this would be reflected in the data by how the serial correlation in shares varies price levels in the market. The parameters $(\theta_1, \theta_2, \theta_3)$ are identified by how these serial correlation patterns covary with demographic characteristics.

The presence of entry aids the identification of the dynamic parameters. In our model, new brands have zero affiliated consumers when they enter the market. The presence of dynamic demand allows these new entrants to initially price lower (as shown in Figure 4) and have low shares, and, over time, increase both prices and shares as they accumulate affiliated consumers. The estimate parameters reflect the magnitudes of these patterns. In the 241 markets in our data, 90 experience the entry of a new brand.

4.2 Implementation

Reduced Computational Complexity

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 computation 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.

Above, 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

$$\sum_z r_{zt} \cdot s_{0t}(z) \exp(\sigma_{jt}(z)) = \sum_z r_{zt} s_{0t}(z) + (\exp(\sigma_{jt}(j)) - 1) r_{jt} s_{0t}(j). \quad (24)$$

By substituting this expression into equation (18), we can obtain a quadratic equation for each

of the $\{s_{0t}(j)\}$:

$$0 = \lambda \frac{1}{s_{0t}(0)} r_{jt} (\exp(\sigma_{jt}(j)) - 1) s_{0t}(j)^2 + \left([\exp(\sigma_{jt}(j)) - 1] (S_{jt} - \lambda r_{jt}) + \lambda \frac{1}{s_{0t}(0)} \sum_{z \in 0, J} r_{zt} s_{0t}(z) + (1 - \lambda) \right) s_{0t}(j) - \lambda \sum_{z \in 0, J} r_{zt} s_{0t}(z) - (1 - \lambda) s_{0t}(0)$$

Conditional on the dynamic parameters and observables, there are only two remaining unknowns: $s_{0t}(0)$ and $\sum_{z \in 0, J} r_{zt} s_{0t}(z)$. 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.

Estimation Routine

To implement our estimator, we use a nested regression approach with the following steps:

1. First, pick the unobserved heterogeneity parameters corresponding to λ_m and $\sigma_{jmt}(j)$.
2. Starting with the first period, solve for $s_{0mt}(0)$ and $\sum_{0,z} r_{zt} s_{0mt}(z)$ in each market using the non-linear system of equations obtained previously. Using the initial value r_{jmt}^* for the first period, iterate forward for each period using the calculated value of r_{jmt} for subsequent periods.
3. Solve for $s_{jmt}(0)$ for each firm, period, and market.
4. Run the regression implied by equation (14) using the $\{s_{jmt}(0)\}$ obtained in the previous step to solve for the linear parameters (α, π, γ) . Calculate the correlation of the residuals $Corr(\hat{\eta}_{jmt}, \hat{\eta}_{jmt(t+1)})$ within each market.
5. Repeat 1-4 to find the unobserved heterogeneity parameters that minimize the sum of squared correlations.

The regression for equation (14) may involve instrumental variables and the use of panel data methods such as fixed effects. In our empirical application, we make use of both.

The estimation methodology employs two tricks to speed up the computation of the dynamic model. First, the explicit formula for $\{s_{jt}(0)\}$ means that the non-linear solver²⁸ only has to find two parameters, $s_{0t}(0)$ and $\sum_{0,z} r_{zt} s_{0t}(z)$, for each market-period. The quadratic form

²⁸To solve for these unknowns, we use a modified contraction mapping that uses the average of the previous guess and the implied solution for the two parameters in each market. This modification improves stability.

Table 3: Summary Statistics by County

Statistic	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max	N
Num. Brands	4.52	1.46	1	3	6	8	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	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.

for the remaining unknowns results in fast calculation. Second, the linear form for the nested regression allows for a quick calculation of the inner part of the routine and allows for serial correlation in unobservables.

In models with state dependence and unobserved heterogeneity, one consideration is the initial value of the unobserved state. Because we have a very long panel, with 104 separate time observations, we have a sufficient “burn-in” period where this issue does not materially affect our estimates.²⁹ As a baseline, we set r_{jmt}^* equal to the value of the aggregate share S_{jm0} in the period prior to the first week of our sample. This value would be consistent with the steady state with no affiliation shock ($\bar{\xi} = 0$). Because we find $\bar{\xi} > 0$, we typically find that $r_{jmt} > S_{jm(t-1)}$ in our estimated model.

4.3 Data for Structural Model

We supplement the EIA-adjusted weekly brand-county share measures with the average prices for the brand in a week-county. To reduce the occurrence of zero shares, which do not arise in the logit model, we use a simple linear interpolation for gaps up to two weeks. We assume that the station was not in the choice set when shares are missing or zero.³⁰ We drop any observations that have missing prices, missing shares, or missing shares in the previous week. This includes dropping the first week of data, for which we do not have previous shares.

Table 3 provides summary statistics of the data for the 241 counties in KY and VA. There is cross-sectional variation in wholesale prices, margins, and the number of stations in each county. To reduce the sensitivity of the analysis to brands with small shares and to make the counterfactual exercises more computationally tractable, we aggregate brands with small shares into a synthetic “fringe” brand. We designate a brand as part of the fringe if it does not appear in ten or more of the 241 markets (counties). Additionally, if a brand does not make up more than 2 percent of the average shares within a market, or 10 percent of the shares for the periods in which it is present, we also designate the brand as a fringe participant for that

²⁹We find that it takes approximately 6 weeks for these values to converge to the steady state, i.e., the choice of the initial value does not substantially affect periods 7 through 104.

³⁰After all cleaning steps, approximately 1 percent of observations have zero shares, and over 90 percent of these zero share occurrences are in spells of 6 weeks or longer.

market. This reduces the maximum number of brands we observe in a county to 8, down from 24. Across all markets, we analyze the pricing behavior of 16 brands, including the synthetic fringe.³¹

We also take steps to reduce measurement error in the number of stations in our data. We assume that stations exist for any gaps in our station-specific data lasting less 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 our sample. After cleaning, we retain 110,844 observations in our sample.

4.4 Results: Demand Estimation

For the empirical application, we implement the methodology described in Section (4.1). Conditional on dynamic parameters, we extract the unobserved shares for all unaffiliated type consumers, obtaining the baseline utility δ_{jmt} . We then estimate static demand parameters using instrumental variables regression following equation (21). Our chosen dynamic parameters minimize the average correlation in brand-market shocks over time (between contemporaneous and a single-period lag), where the correlation is calculated within each market.³²

Our regression equation takes the following form:

$$\ln \left(\frac{s_{jmt}(0)}{s_{0mt}(0)} \right) = \alpha p_{jmt} + \pi (p_{jmt} \times Income_m) + \gamma N_{jmt} + \zeta_{jm} + \phi_t + \psi_{m,month(t)} + \eta_{jmt}$$

Here, the subscript m denotes the market (county). We have shares and prices at the brand-county-week level. Within-county shares of unaffiliated consumers depend on prices, station amenities, and demographic characteristics of the local population. The brand-county fixed effects, ζ_{jm} , control for variation in the number of stations, brand amenities, and local demographic characteristics. Because we observe station entry and exit, we also include the number of stations for the brand in that market, N_{jmt} , to capture within-brand-county variation over time.³³

Thus, brand-county fixed effects, which are identified by the panel, allow us to account for a first-order component of heterogeneity in preferences. Another important component of preferences in this model is price sensitivity. To account for heterogeneity in price sensitivity, we interact price with the log median household income in the county.³⁴

³¹Summary statistics by brand are presented in Table 16 in the Appendix. The fringe brand is, on average, 13 percent of the shares for the markets that it appears in. As we designate a fringe participant in nearly every market, the aggregated fringe has the highest overall share (12 percent).

³²In the estimated model, the mean market-level correlation in shocks is -0.01.

³³Station amenities include, for example, the presence of food (snack or restaurant), co-location with a super-market, car services, and proximity to an interstate. Demographic characteristics might include median household income, population, population density, and commute percent. These do not vary much over time in our sample. We do not directly account for commuting patterns in our analysis.

³⁴We do not control for unobserved heterogeneity in price sensitivity, which would add a significant computational

In addition to the brand-county fixed effects, we employ panel data methods to address other unobservables. We allow for the fact that δ_{jmt} may be correlated over time in ways not dependent on (p, N, ζ) . We let the time-varying unobserved components of demand be specified as $\phi_t + \psi_{m,month(t)} + \eta_{jmt}$. That is, we estimate period (weekly) fixed effects $\{\phi_t\}$ and county-specific (monthly) seasonal demand shocks $\{\psi_{m,month(t)}\}$.³⁵ Once we incorporate these fixed effects, the identifying restriction for the dynamic parameters is that the brand-market-period specific shock η_{jmt} is uncorrelated across periods, after accounting for aggregate period-specific shocks, county-level seasonal patterns, and brand-county level differences. Thus, our model attributes the residual brand-specific correlation in demand over time within a market to unobservable consumer types arising from affiliation.

We allow for endogeneity in pricing behavior by instrumenting for p_{jmt} with deviations in wholesale costs arising from crude oil production in the US. The instrument (z_1) is constructed from a regression of deviations of wholesale costs (from the brand-county average) on the interaction of weekly US production of crude oil with the average wholesale cost for the brand in the county.³⁶ This gives us brand-county-specific time variation in our instrument which is (a) correlated with the wholesale cost and (b) plausibly not linked to demand. We chose this measure, rather than instrumenting directly with brand-state wholesale costs, to allow for the possibility that local variation in wholesale costs over time may reflect brand-specific demand shocks.

We interact the above instrument with $Income_m$ to create a second instrument, z_2 , to account for the endogeneity of $(p_{jmt} \times Income_m)$. Both US crude oil production and income are plausibly exogenous with respect to local, time-varying demand shocks. Figure 7 in the Appendix summarizes the time-series variation by plotting mean total market shares and mean prices during our sample in panel (a). In panel (b), we plot the mean instrument z_1 against the mean price. As the figure shows, there is a strong correlation with the instrument, constructed from US production of crude oil, and prices. Prices display seasonal patterns, reflecting demand, while our instrument does not.

The estimates for the linear parameters are reported in Table 4. The first three columns report coefficient estimates from a logit demand regression using observed shares. The fourth column reports the results for unaffiliated customers from our dynamic model. In the static model, all consumers are assumed to be unaffiliated. We obtain a larger (in magnitude) price coefficient with our dynamic specification, as the model has isolated unaffiliated consumers (including shoppers) from affiliated consumers with lower price sensitivities. As expected,

burden. Despite a wide range in market-level median income in our data (from \$20,000 to \$124,000, or a spread of 1.8 log points) we find modest effects of income on price sensitivity.

³⁵We benefit from the size of our dataset. 95 percent of county-months have at least 18 observations, and 99 percent of county-brands have at least 40 observations.

³⁶Our measure of the average brand-county wholesale cost is the fixed effect obtained by a regression of wholesale costs on brand-county and weekly fixed effects, thereby accounting for compositional differences across time. We obtain the quantity of crude oil produced in the U.S. each week from the EIA.

Table 4: Estimates of Static Demand Parameters

	Static Model			Dynamic Model
	(1)	(2)	(3)	(4)
Price	-0.022* (0.014)	-0.260*** (0.038)	-2.315*** (0.507)	-7.615*** (1.862)
Price \times Income	-0.135*** (0.021)	0.077*** (0.023)	-0.007 (0.032)	0.329*** (0.129)
Number of Stations	0.016*** (0.004)	0.064*** (0.010)	0.063*** (0.010)	0.241*** (0.040)
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, where the dependent variable is the log ratio of the share of the brand to the share of the outside good. 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, which depend on the estimated dynamic parameters. Standard errors are clustered at the county level. For the dynamic model, standard errors are calculated via the bootstrap.

higher income corresponds to a lower price sensitivity, though this effect is small. A 1 log point (or 271 percent) increase in income changes the price coefficient from -7.615 to -7.285 . We find that an increase in the number of stations that a brand has in a market has a statistically significant positive effect on demand for unaffiliated consumers.

Table 5 reports estimates of the dynamic parameters. The parameters θ_1 , θ_2 , and θ_3 imply that 62 percent of consumers, on average, are subject to state dependence and develop an affiliation for the brand they previously purchased from. The coefficient of 0.049 on income indicates that higher-income consumers are more likely to develop an affiliation, consistent with a habit-formation model where switching costs are increasing in wages. Likewise, the coefficient of 0.019 on population density indicates that consumers in more dense areas are more likely to become affiliated. We interpret this to reflect that urban environments have higher driving costs, which also increase switching costs. Both of these demographic variables are standardized, so each coefficient corresponds to an increase of one standard deviation. Neither is statistically significant at the 95 percent level.

The estimated utility shock to affiliated consumers, $\bar{\xi}$, implies that the affiliated consumers are inelastic with respect to price. Across observations, the mean own-price elasticity for affiliated consumers is -0.55 , and the median is zero. The utility shock generates behavior so that affiliated consumers do not respond to prices when the levels are low, but they do when the

Table 5: Estimates of Dynamic Demand Parameters

	<i>Baseline</i> θ_1	Affiliation Rate <i>Income</i> θ_2	<i>Density</i> θ_3	Strength of Affiliation <i>Utility Shock</i> $\bar{\xi}$
Coefficient	0.479	0.049	0.019	18.682
95 Percent CI	[0.45, 0.50]	[-0.05, 0.08]	[-0.01, 0.04]	[18.68, 19.06]

Notes: Table displays the estimated non-linear coefficients from the dynamic model. The first three parameters imply that, on average, 62 percent of consumers that purchase develop an affiliation for that brand. Brands located in areas with higher incomes and higher population densities have greater rates of affiliation, though this heterogeneity is not statistically significant. The last parameter shows the level shock for affiliated customers, which is positive, as expected. Confidence intervals are shown with brackets and are calculated via the bootstrap.

prices are high enough. The average (absolute) weekly price change in the data is 5 cents per gallon, and the 25th and 75th percentile price changes are 1.6 cents and 7.5 cents, respectively. We therefore find that affiliated consumers do not typically switch brands within this range of price changes.³⁷ The unaffiliated consumers, however, are highly elastic, with an average own-price elasticity of -21.1 . This is large in magnitude, and it implies that for a 1 percent increase in price (roughly 3 cents), the station will lose 21.1 percent of the unaffiliated consumers. This high level of price sensitivity for a subset of retail gasoline consumers seems plausible, as some “shoppers” have been found to go well out of the way to save a few cents per gallon.³⁸

On average, roughly 26 percent of a brand’s customers come from unaffiliated consumers in any week in equilibrium. The average weighted elasticity, which weighs affiliated and unaffiliated consumers by their relative (purchasing) proportions, is -3.4 . This weighted elasticity captures the effective elasticity faced by a firm and is different than the elasticity obtained when estimating a static model. A “naive” estimate using a static model would result in a value of -5.7 , which implies a much greater loss in market share for a given price increase than we estimate from the dynamic model. Our brand-county elasticity estimate of -3.4 is more inelastic than some other estimates in the literature, which may be due to the fact that the existing estimates do not account for consumer inertia.³⁹ At the market level, the parameter estimates imply an aggregate weekly elasticity of demand of -2.5 .⁴⁰

³⁷A common limitation of empirical models is that it is only possible to capture local variation in the data. Our model implies that that affiliated consumers become more elastic at higher price levels. It is likely these consumers would become even more elastic if subject to price increases of a much greater magnitude. Thus, firms end up pricing to both affiliated and unaffiliated consumers in equilibrium.

³⁸For example, the National Association of Convenience Stores found in their 2018 survey that 38 percent of people would drive 10 minutes out of their way to save 5 cents per gallon. See, <https://www.convenience.org/Topics/Fuels/Documents/How-Consumers-React-to-Gas-Prices.pdf>

³⁹Additionally, the estimates of Houde (2012) of -10 to -15 reflect an elasticity at the station level, rather than at the brand-county level, which should result in more elastic estimates.

⁴⁰The literature has typically estimated the aggregate market elasticity for gasoline to be highly inelastic. For example, Levin et al. (2017) find an aggregate elasticity of -0.30 and Li et al. (2014) estimate it to be -0.1 . An important distinction between our estimate and previous aggregate elasticities is that ours is obtained using more

We have tested the robustness of our demand estimates along two important dimensions. First, we consider the impact of the initial value of the unobserved state on our estimates. As discussed in Section 4.2, because our dataset consists of a long panel with 104 time periods, the initial choice of r_{jm0}^* may not matter much in estimation. After several periods, the system converges toward the steady-state regardless of the initial value. The estimates will be less sensitive to the initial choice if more observations reflect the steady-state values. To evaluate the potential impact, we consider two extreme choices of r_{jm0}^* . First, we assume that no affiliation exists at the start of the sample, imposing $r_{jm0}^* = 0$ for all brands and markets. At the other extreme, we choose r_{jm0}^* to reflect a scenario where all consumers are affiliated in the first period, including shoppers. Each extreme scenario delivers similar point estimates that fall within the 95 percent confidence intervals of our baseline estimates.

Second, we test for the uniqueness of our solution by randomly drawing initial values for the dynamic parameters in the estimation algorithm. We confirm that our estimation routine gets very similar results from most starting values. In particular, the estimation routine converges to nearly-identical estimates of $\bar{\xi}$ and θ_1 regardless of the starting value. The estimates for θ_2 and θ_3 remain close to zero, and the variation in these parameters due to starting values has no appreciable impact on the value of λ .

5 Empirical Application: Supply-Side Analysis

5.1 Dynamic Pricing Behavior

Given the demand estimates, we construct the components in each firm's Bellman equation from (6). We assume that each brand sets prices to maximize discounted profits at the county level. The dynamic condition for optimal pricing for brand j is:

$$\frac{\partial \pi_{jt}}{\partial p_{jt}} + \beta \frac{\partial E [V_j(r_{t+1}, c_{t+1}, x_{t+1}) | p_t, r_t, c_t, x_t]}{\partial p_{jt}} = 0, \quad (25)$$

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.

The estimation of dynamic parameters, along with our measures of marginal costs, allow for a direct estimate of the derivative of the static profit with respect to price: $\frac{\partial \pi_{jt}}{\partial p_{jt}}$. If this were zero, it would imply that firms are pricing myopically in the context of the model, as they are simply maximizing the current-period profits. When it is non-zero, it implies that dynamic considerations are affecting a firm's pricing decision.

narrowly defined markets and over a shorter time horizon (one week). Both of these features should produce more elastic estimates. Levin et al. (2017) do estimate a two-day (aggregate) elasticity of -1.38 , which is closer to our estimate.

Table 6: Summary of Implied $\beta \frac{\partial E[V_j(\cdot)]}{\partial p_{jt}}$

Group	Mean	Min	p25	Median	p75	Max
All	-0.109	-0.761	-0.152	-0.083	-0.042	0.628

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.

On average, we find that $\frac{\partial \pi_{jt}}{\partial p_{jt}}$ is positive. This implies that firms are systematically pricing lower than the myopic profit-maximizing price. We interpret this as evidence of forward-looking behavior and the presence of dynamics, consistent with the reduced-form evidence of Section 3.3.2. Based on equation (25), we attribute the difference between $\sum_{l \in J_i} \frac{\partial \pi_{lt}}{\partial p_{jt}}$ and 0 to be accounted for by the derivative of the continuation value (DCV), $\beta \frac{\partial E[V_j(\cdot)]}{\partial p_{jt}}$. That is, 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.⁴¹ After estimating demand in an independent step, we are able to recover these residuals directly.

Summary statistics for the value of the derivative of the continuation value (DCV) are presented in Table 6. 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 significant: the mean of -0.109 implies that a 1 cent increase in price would increase static profits by roughly 4 percent.⁴² Intuitively, firms are lowering prices to invest in future demand. Such behavior allows firms to occasionally have negative price-cost 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 Supply-Side Estimation

To estimate counterfactual pricing behavior by firms, it is necessary to estimate how dynamic incentives vary with state variables and firm actions. Two approaches are possible. The first is to take a stance on the beliefs of firms and, via forward simulation, solve for the equilibrium continuation value function. We develop an alternative approach in which we rely on the structural demand model to calculate the static component of profits, and we use a functional approximation to capture dynamic incentives.

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. This approach

⁴¹Other explanations may be plausible. For example, a component of this residual may be profits obtained by complementary products, such as food sold at retail gasoline stations.

⁴²The average (scaled) profit in our data is 0.029. Price-cost margins are approximately 21 cents per gallon.

has three key advantages. First, the static component of profits is obtained without having to make any assumptions about firm expectations and discount rates. Second, utilizing an approximation for the DCV greatly reduces the computational time needed to re-compute equilibria. Third, we avoid the need to make dimension-reducing assumptions, such as constructing a limited grid for prices, that are less palatable in our setting. To accurately represent behavior, this approach requires 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.

Using the data and the estimated demand parameters, we obtain estimates of the DCV and project these estimates on prices and state variables, including measures that capture expectations. In general, Markovian assumptions allow for the continuation value to be expressed as a function of the state variables and actions of the firms. We estimate the following dynamic first-order condition:

$$\frac{\partial \pi_{jt}}{\partial p_{jt}} + \Psi_j(p_t, r_t, c_t, x_t; \theta) + \zeta_{jt} = 0. \quad (26)$$

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 $\beta \frac{\partial E[V_j(\cdot)]}{\partial p_{jt}}$ from equation (25), 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. 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. One advantage of our approach is that it is not necessary to take a stance on the discount rate or the beliefs of firms; both of these are absorbed into the reduced-form model.

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 the future profits, consistent with the approximation used in estimation. Thus, our approach can be considered as an attempt to replicate the forecasting behavior of firms that use regression analysis to predict future profits. In this 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 (26). In other words, we assume that firms' forecasts of the DCV correspond to $\Psi(\cdot)$, and we evaluate how close these forecasts are to the realized DCV when firms choose prices following these forecasts. We discuss these forward simulations below.

To estimate $\Psi_j(\cdot)$, we project the estimated DCV onto the share of affiliated consumers r_{jmt} , marginal costs, and expectations of future costs. Our model and descriptive evidence suggests

Table 7: Dynamic Pricing Incentive: Regressions

	$\beta \frac{\partial E[V_j(\cdot)]}{\partial p_{jt}}$	Sensitivity
	(1)	(2)
r_{jmt}	-0.625*** (0.001)	5.864*** (0.017)
$r_{jmt} \times \frac{\partial S_{jmt}}{\partial p_{jmt}}$	-0.147*** (0.002)	2.329*** (0.044)
Marginal Cost	-0.003*** (0.000)	0.042*** (0.004)
Cost Change (30-Day Ahead)	0.019*** (0.001)	-0.203*** (0.019)
λ_m	-0.181*** (0.012)	3.070*** (0.268)
Num. Stations (Brand)	-0.000*** (0.000)	0.010*** (0.000)
Num. Stations (Market)	-0.000*** (0.000)	0.000** (0.000)
Num. Brands (Market)	-0.002*** (0.000)	0.051*** (0.002)
Constant	Yes	Yes
Observations	110,844	110,844
R ²	0.912	0.625

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.

that these variable play an important role in expectations of future profits. We also include the interaction of r_{jmt} with the derivative of shares with respect to price, motivated by the fact that future considerations may be directly influenced by the derivative of current period shares. 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 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 (26) are reported in Table 7. The first specification reports the coefficients from a regression of the DCV onto the dependent variables. 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. The nine-parameter model has an R^2 of 0.91. Overall, the reduced-form approach captures the vast majority of the price variation that cannot be explained by static optimization.⁴³

We find that a higher share of affiliated consumers at the firm level r_{jmt} and the market level λ_m increases the magnitude of the DCV, corresponding to an increased investment incentive when pricing. These variables have the largest impact on the DCV. The effect of affiliated consumers on pricing behavior is exacerbated when the consumers are more responsive to price in the current period, as indicated by the coefficient of -0.147 on $r_{jmt} \times \frac{\partial S_{jmt}}{\partial p_{jmt}}$. Consistent with forward-looking behavior, increases in marginal costs today, which tend to reduce current profits, make firms more sensitive to dynamic profit considerations. Conversely, expected increases in future costs lead firms to place relatively more weight on current profits. Relative to the impact of affiliated consumers, we find that other market-level controls have modest effects.

To help interpret how sensitive firms are to dynamic considerations, the second column of Table 7 reports a regression where we replace the value of the DCV with the logged absolute value. Thus, the 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. Typically, a negative coefficient in the first column corresponds to a positive coefficient in the second, as the average value for the DCV is negative.

Verification with Forward Simulations

To verify that our estimate of $\Psi(\cdot)$ is consistent with realized profits, we use forward simulations to calculate the discounted present value of per-period profits when a firm unilaterally deviates its price. For each brand in each market and each period, we slightly perturb the price for that brand. We re-compute shares in that period and then calculate equilibrium play in future periods, while imposing that firms construct beliefs using our estimate of $\Psi(\cdot)$. We use the incremental change in profits along this simulated path to evaluate $\beta \frac{\partial E[V_j(\cdot)]}{\partial p_{jt}}$ under rational expectations. Using an annual discount factor of 0.9, we calculate an average simulated value for the derivative of -0.119 , which compares favorably to the mean estimate of -0.109 in Table 6. Thus, our approach to capturing firm expectations seems reasonably well-aligned with realized equilibrium profits. For additional details about the simulations, see Appendix E.

⁴³The high degree of explanatory power of the parsimonious model provides some confidence for reasonable interpolation and extrapolation in counterfactual analysis.

5.3 Horizontal Market Power: Merger Simulation

To evaluate the impact of dynamic pricing incentives on horizontal market power, we simulate a merger between Marathon and BP, which are the number one and number four (non-fringe) brands in terms of overall shares in our sample. Out of the 241 markets, they overlap in 75. In these 75 markets, the average (post-merger) HHI is 1511, and the mean change in HHI resulting from the merger is 383. In 8 markets, the resulting HHIs are greater than 2500, and the changes are greater than 200, meeting the typical thresholds that are presumed likely to enhance market power. The merger would change twelve markets from 3 firms to 2 firms and eighteen markets from 4 firms to 3 firms. We allow the firms to merge at the beginning of September 2014, and we calculate counterfactual prices and shares for the second half of the sample.⁴⁴

In Section 2, we showed that the price effects of a merger can depend on the way the merger is implemented, especially in the presence of consumer inertia. To measure the potential empirical impact, we consider the two alternative merger scenarios described in that section: a joint pricing merger and a brand consolidation merger. In the joint pricing scenario, the merged firm has pricing control over both brands, which they maintain as distinct entities. In the brand consolidation scenario, the merged firm consolidates the assets under a single brand. To implement a brand consolidation merger, we need to make an assumption about how much utility consumers of the product removed from the market will receive from buying the consolidated brand. As in the theoretical analysis, we assume that, at pre-merger prices, the consolidated brand will have the same combined share of unaffiliated customers as the separate pre-merger brands.

For the joint pricing scenario, we also need to make an assumption about the cross-price effects on the continuation value, i.e., $\beta \partial E[V_k] / \partial p_{jt}$ when k and j are owned by the merged firm. For our baseline results, we assume that the effects are proportional to the diversion ratios D_{kj} , so that $\partial E[V_k] / \partial p_{jt} = D_{jk} \partial E[V_k] / \partial p_{kt}$. The diversion ratios capture the relative effects on shares, which should be correlated with the effects on profits. Our counterfactual results are qualitatively similar with moderate changes in this scaling factor, such as assuming the cross-price effects are zero. As expected, if we reduce the impact of the cross-price effects on the continuation value, we get a lower impact on post-merger prices.

Table 8 displays the mean effects on prices, shares, and profits of the two mergers. The first three columns report the effects from the joint pricing scenario, and the second three columns report the effects from the brand consolidation scenario. In either scenario, the overall price effects are moderate. The joint pricing scenario predicts that the merging firms will raise prices by 3.1 percent, with an 11 percent decrease in shares and a meaningful increase in profits. Rival

⁴⁴Because 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 that exceed this value. The baseline functional form of demand may not be reasonable for extreme out-of-sample values. Across all merger counterfactuals, only 8 observations approach the choke price.

Table 8: 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	3.14	-11.07	20.12	1.97	0.11	30.43	5.35	-16.74	34.44
Other	0.05	4.64	5.29	1.82	-3.59	17.21	0.61	3.86	11.28
Overall	1.26	-2.42	12.4	1.90	-1.93	23.55	2.40	-5.40	22.39

Notes: Table displays the mean effects 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 second three columns report a counterfactual brand consolidation merger. The last three columns provide the estimates for a brand consolidation merger from a static model that is calibrated to match prices, margins, and shares from the same data. Price effects are weighted by share.

firms in the same market see an increase in share and a small effect on prices, which results in higher profits. Overall, prices in the market increase by 1.3 percent.

The brand consolidation scenario predicts a price increase of 2.0 percent, which is lower than the joint pricing prediction. However, in the brand consolidation scenario, the merging firms realize a slight *increase* in market share. This occurs because of asymmetries among firms. When merging firms have different brand values, the affiliated consumers of the weaker brand become more attached to the consolidated brand post-merger. In other words, affiliated consumers benefit from becoming locked into a better brand. Recall that our merger holds the choice probabilities for unaffiliated consumers to be identical post-merger at the pre-merger prices. In this scenario, if we hold prices fixed at the pre-merger levels, we would find that the merging firm accumulates greater shares over time due to the superior strength of affiliation.

Another difference between the two mergers is in the effects on rivals. In contrast to the joint pricing scenario, the brand consolidation scenario sees rivals increase prices by 1.8 percent, which is almost as much as the merging firms. This leads to a greater overall price increase of 1.9 percent across the affected markets and higher firm profits.

Overall, these effects are modest relative to what might be expected from a merger with the corresponding market shares and HHIs. For comparison, we report the results from a merger analysis using a static model in the last three columns of Table 8. We calibrate a standard logit demand system to identical prices, margins, and shares that are used to estimate the dynamic model. We use a brand consolidation merger to illustrate the potential effects.⁴⁵ The static model predicts price effects of over 5 percent for the merging firms, which is meaningfully higher than the predictions from the dynamic models.

Consistent with our simulations in Section 2, we find that the dynamic model with consumer

⁴⁵Though we showed in Section 2 that both types of mergers predict identical price effects in a symmetric oligopoly, they can have different predictions with asymmetries.

Table 9: Equilibrium Effects of State Dependence

Brand	Increase λ by 0.07			Decrease λ by 0.07		
	Price	Share	Profit	Price	Share	Profit
Marathon-BP	2.03	6.66	33.81	-1.73	-7.82	-28.5
Other	1.40	7.99	26.44	-1.64	-8.24	-26.27
Overall	1.68	7.39	29.97	-1.68	-8.05	-27.34

Notes: Table displays the mean effects from a counterfactual scenarios with different portions of consumers affected by state dependence. The first three columns report the equilibrium effects with a greater prevalence of state dependence, and the second three columns report the effects with less state dependence.

inertia can predict smaller price increases than 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.

5.4 Dynamic Market Power

In the previous section, we used the empirical model to evaluate the role of horizontal market power in the presence of consumer inertia. In this section, we isolate the role of dynamic market power. To do so, we change the share of consumers that are affected by inertia, λ .

We simulate counterfactual scenarios where the market-specific λ increases or decreases by 0.07. These changes correspond approximately to an 11 percent change in λ , as the average value across markets is 0.62. Thus, we consider the impact of the share of “shoppers” on equilibrium prices. Though we do not consider a merger in this section, we report results for Marathon and BP separately from other brands and only for the markets where both firms are present. This facilitates a comparison to the horizontal market power effects in Table 8.

Table 9 reports the results from the counterfactual scenarios. The first three columns report the equilibrium effects of an increase in the prevalence of state dependence, and the second three columns report the effects of a decrease. If an additional 7 percent of consumers in the market were affected by state dependence, then overall prices would increase by 1.7 percent. Marathon and BP would benefit slightly more than other firms, increasing prices by 2 percent and realizing a 34 percent increase in profits. The effects of reducing λ by the same amount are quite similar, but in the opposite direction.

Overall, these effect on prices and profits are similar to the effects of a brand consolidation merger between Marathon and BP (Table 8). In rough terms, reducing the share of consumers that are shoppers by 7 percentage points has equilibrium price effects on par with the effects of a 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 switching costs (or consumer inertia, more generally) may

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

6 Conclusion

We develop a model of consumer inertia that accounts for commonly observed dynamic pricing behavior, such as the slow adjustment of prices to changes in cost. The dynamics result from competing firms optimally setting prices to consumers that may become loyal or habituated to their current supplier. Dynamic market power reflects the ability of firms to raise price in response to consumer inertia relative to static consumer demand, though consumer inertia can lead to lower prices in equilibrium, depending upon whether the incentive to harvest affiliated consumers dominates the incentive to invest in future demand. High levels of dynamic market power may correspond with lower levels of horizontal market power, i.e., the impact of competition from rival products on prices.

Using data from retail gasoline markets, we first present reduced-form evidence consistent with consumer inertia. Firm-level shares are highly correlated over time, prices slowly adjust to changes in marginal costs, and prices anticipate future changes in expected costs. The evidence suggests that, even in a relatively competitive market with a homogeneous product, accounting for dynamics in consumer behavior and firm behavior may be important.

We develop an empirical model that can identify dynamic demand parameters using data on price, shares, and an instrument. Results suggest that 62 percent of retail gasoline consumers become affiliated to the firm from which they currently purchase on a week-to-week basis, and that these consumers are price insensitive. Conversely, we find that unaffiliated consumers are quite price sensitive and play an important role in disciplining equilibrium prices. We evaluate the dynamic incentives that affect prices. We show, both theoretically and empirically, that failing to account for dynamic demand can cause significant biases when predicting post-merger price increases. Counterfactual analyses show that a relatively small fraction of consumers, through dynamic market power, can impact prices as much as a change in competition arising from a merger.

References

- AGUIRREGABIRIA, V. AND P. MIRA (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, 75, 1–53.
- ARCIDIACONO, P. AND R. A. MILLER (2011): “Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity,” *Econometrica*, 79, 1823–1867.
- 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. VALIMAKI (1996): “Learning and Strategic Pricing,” *Econometrica*, 64, 1125.
- BERGEMANN, D. AND J. VÄLIMÄKI (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,” Working Paper w23829, National Bureau of Economic Research, Cambridge, MA.
- 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.
- 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.
- 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.
- ECKERT, A. (2013): “Empirical Studies of Gasoline Retailing: A Guide to the Literature,” *Journal of Economic Surveys*, 27, 140–166.

- 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.
- ENGLE, R. F. AND C. W. J. GRANGER (1987): “Co-Integration and Error Correction: Representation, Estimation, and Testing,” *Econometrica*, 55, 251.
- FARIAS, V., D. SAURE, AND G. Y. WEINTRAUB (2012): “An Approximate Dynamic Programming Approach to Solving Dynamic Oligopoly Models,” *The RAND Journal of Economics*, 43, 253–282.
- 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.
- 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.
- HECKMAN, J. J. (1981): “Heterogeneity and State Dependence,” *Studies in Labor Markets*, 91–140.
- HENDEL, I. AND A. NEVO (2013): “Intertemporal Price Discrimination in Storable Goods Markets,” *American Economic Review*, 103, 2722–2751.
- 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.
- HOUE, J.-F. (2012): “Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline,” *American Economic Review*, 102, 2147–2182.
- 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.

- LI, S., J. LINN, AND E. MUEHLEGGGER (2014): "Gasoline Taxes and Consumer Behavior," *American Economic Journal: Economic Policy*, 6, 302–342.
- 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, C. RYAN, AND G. SHEU (2016): "Pass-Through and the Prediction of Merger Price Effects: Pass-Through in Merger Analysis," *The Journal of Industrial Economics*, 64, 683–709.
- MILLER, N. H., M. REMER, AND G. SHEU (2013): "Using Cost Pass-through to Calibrate Demand," *Economics Letters*, 118, 451–454.
- NOSAL, K. (2011): "Estimating Switching Costs for Medicare Advantage Plans," Working paper.
- 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.
- SLADE, M. E. (1998): "Optimal Pricing with Costly Adjustment: Evidence from Retail-Grocery Prices," *The Review of Economic Studies*, 65, 87–107.
- STAHL, D. O. (1989): "Oligopolistic Pricing with Sequential Consumer Search," *The American Economic Review*, 700–712.
- SWEETING, A. (2013): "Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry," *Econometrica*, 81, 1763–1803.
- VARIAN, H. R. (1980): "A Model of Sales," *American Economic Review*, 70, 651–659.
- WEINTRAUB, G. Y., C. L. BENKARD, AND B. VAN ROY (2008): "Markov Perfect Industry Dynamics With Many Firms," *Econometrica*, 76, 1375–1411.
- YEO, J. AND D. P. MILLER (2018): "Estimating Switching Costs with Market Share Data: An Application to Medicare Part D," *International Journal of Industrial Organization*, 61, 459–501.

Appendix

A More Details on Theoretical and Numerical Steady-State Analysis

A.1 Monopoly

We analyze steady-state prices in a monopoly market (with an outside good) to show how habit-forming consumers affect optimal prices and markups.

To simplify notation in the monopoly case, let the monopolist's share of affiliated and unaffiliated consumers be s_j and s_0 , respectively, and its number of affiliated consumers be r . Consumers become affiliated at rate λ . We assume positive dependence in purchase behavior, so that $s_j > s_0$. In the steady state, $r_{jt} = r_{j(t+1)} = r_j$ and $c_{jt} = c_{j(t+1)} = c_j$. The steady-state number of affiliated consumers, r^{ss} , is:

$$\begin{aligned} r^{ss} &= \lambda((1 - r^{ss})s_0 + r^{ss} \cdot s_j) \\ \implies r^{ss} &= \frac{\lambda s_0}{1 - \lambda(s_j - s_0)}. \end{aligned}$$

The steady-state number of affiliated consumers is increasing in the probability of becoming affiliated, λ , and the difference between the choice probabilities of affiliated consumers and unaffiliated consumers, $s_j - s_0$. Using the steady-state value of affiliated consumers, we can solve for the steady-state pricing function.

The steady-state period value is:

$$\begin{aligned} V^{ss}(r^{ss}, c^{ss}) &= (p^{ss} - c^{ss})((1 - r^{ss})s_0 + r^{ss}s_j) + \beta V^{ss} \\ &= \frac{p^{ss} - c^{ss}}{1 - \beta} \cdot \frac{s_0}{1 - \lambda(s_j - s_0)}. \end{aligned}$$

This equation represents the monopolists 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 $s_j = s_0$), in which case profits are simply the per-unit discounted profits multiplied by the firm's market share.

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

$$p^{ss} = c^{ss} + \underbrace{\frac{-s_0(1 - \lambda s_j + \lambda s_0)}{\frac{ds_0}{dp}(1 - \lambda s_j) + \frac{ds_j}{dp}\lambda s_0}}_{m = \text{markup of price over marginal cost}}. \quad (27)$$

The second term, m , on the right-hand side of equation (27) 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 (27) 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:

$$m_d = -\frac{s_0(1 - \lambda s_j + \lambda s_0)}{\frac{ds_0}{dp}(1 - \lambda s_j) + \frac{ds_j}{dp}\lambda s_0} \begin{matrix} \geq \\ \leq \end{matrix} -\frac{s_0}{ds_0/dp} = m_s.$$

For a given price, the terms s_0 and ds_0/dp are equivalent across the two models. Rearranging terms, we obtain a simple condition relating the levels of the markup terms:

$$m_d > m_s \iff -\frac{\partial s_0}{\partial p} > -\frac{\partial s_j}{\partial p}. \quad (28)$$

A higher value for m_d indicates higher markups and higher prices. Thus, if affiliated consumer quantities are relatively less sensitive to changes in price, then markups are higher.

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 (28). 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 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.

A.2 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)

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

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 $\frac{dp}{dr}$ and repeat steps 1-4 above until a solution is found.

A.3 Simulation: Example Markets

Figure 6: Monopoly and Oligopoly Prices with Consumer Affiliation



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

To illustrate the impact of affiliation on pricing incentives, we plot the equilibrium prices for two different sets of utility parameters in Figure 6. 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 on equilibrium prices, and that the impact on price may be non-monotonic in the proportion of consumers prone to affiliation. The non-monotonicity is a result of two countervailing affects. At low levels of λ firms face increasing 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.

A.4 Numerical Simulation Parameters

Table 10: 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), margins between 0.05 and 0.75, and converged for all values of $\lambda \in \{0.05, 0.1, 0.15, \dots, 0.70\}$. See the text for additional details.

A.5 Outcomes and Demand Parameters

Table 11: Joint Pricing Merger Simulation: Demand Parameters

	(1) Pre-Merger Price	(2) Price Δ F1	(3) Bias F1	(4) Price Δ F3	(5) Bias F3
λ	0.128*** (0.017)	-1.538*** (0.093)	4.190*** (0.060)	0.156*** (0.038)	8.680** (4.287)
α	0.144*** (0.003)	2.553*** (0.017)	0.674*** (0.011)	0.908*** (0.007)	-15.754*** (0.802)
ξ	0.026*** (0.003)	1.908*** (0.017)	0.433*** (0.011)	0.772*** (0.007)	-15.215*** (0.792)
$\bar{\xi}$	0.040*** (0.002)	-0.163*** (0.013)	0.543*** (0.008)	0.186*** (0.005)	1.607*** (0.588)
Constant	1.883*** (0.012)	8.157*** (0.067)	-0.451*** (0.043)	0.971*** (0.028)	29.107*** (3.104)
N	6566	6566	6566	6566	6566

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variables are outcomes from a merger simulation that consolidates pricing control of products 1 and 2. F1 and F3 refer to firms 1 and 3, respectively. Pre-Merger price is for firm 1. Price Δ and Bias are the merger price change and the absolute value of the simulation bias, respectively. Parameters correspond to the dynamic demand model detailed in section 2. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Tables 11 and 12 provide results from regressions of pre-merger prices, price changes, and the absolute value of the bias in the static prediction on the demand parameters. Pre-merger prices increase with the fraction of affiliated customers (λ) and the strength of affiliation ($\bar{\xi}$). For joint pricing mergers, stronger dynamics reduce merger price effects and generate a greater under-prediction that arises from a misspecified static model, which increases the size of the bias. In brand consolidation mergers, stronger dynamics also generate larger merger price changes and absolute bias. However, these relationships do not hold in every instance, and may interact in interesting ways. As we show in Figures 6 and 1, prices are often non-monotonic in λ .

Table 12: Brand Consolidation Merger Simulation: Demand Parameters

	(1) Pre-Merger Price	(2) Price Δ F1	(3) Bias	(4) Price Δ F3	(5) Bias F3
λ	0.128*** (0.017)	3.609*** (0.115)	1.406*** (0.052)	1.144*** (0.039)	-676.899 (973.032)
α	0.144*** (0.003)	2.846*** (0.022)	-0.282*** (0.010)	0.808*** (0.007)	-79.986 (182.082)
ξ	0.026*** (0.003)	1.754*** (0.021)	-0.512*** (0.010)	0.691*** (0.007)	-144.450 (179.863)
$\bar{\xi}$	0.040*** (0.002)	0.830*** (0.016)	0.410*** (0.007)	0.154*** (0.005)	9.927 (133.556)
Constant	1.883*** (0.012)	8.313*** (0.083)	0.802*** (0.038)	0.505*** (0.028)	930.684 (704.385)
N	6566	6566	6566	6566	6566

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variables are outcomes from a merger simulation that consolidates products one and two under one brand. F1 and F3 refer to firms 1 and 3, respectively. Pre-Merger price is for firm 1. Price Δ and Bias are the merger price change and simulation bias, respectively. Parameters correspond to the dynamic demand model detailed in section 2. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.6 Numerical Results for Merging Firms

Table 13: Joint Pricing Simulation: Merger Price Change and Bias

	(1) Price Δ	(2) Price Δ	(3) Price Δ	(4) Bias	(5) Bias	(6) Bias
Pre-Merger Market Share	27.834*** (0.276)	29.122*** (0.248)	10.971*** (1.522)	4.582*** (0.223)	3.264*** (0.185)	0.094 (0.945)
Pre-Merger Margin	12.870*** (0.145)	13.087*** (0.129)	14.481*** (0.261)	7.184*** (0.118)	6.963*** (0.096)	12.024*** (0.196)
λ		-3.612*** (0.087)	-2.774*** (0.110)		3.696*** (0.065)	3.225*** (0.071)
α			1.013*** (0.100)			0.608*** (0.062)
ξ			1.039*** (0.098)			0.763*** (0.061)
$\bar{\xi}$			-0.446*** (0.012)			0.265*** (0.008)
Price Δ						-0.252*** (0.008)
Constant	-4.410*** (0.058)	-3.337*** (0.058)	0.250 (0.276)	-1.224*** (0.047)	-2.322*** (0.043)	-3.583*** (0.171)
N	6566	6566	6566	6566	6566	6566

Notes: The dependent variables are outcomes from a merger simulation that consolidates pricing control of products 1 and 2. Observations are for firm 1. Price Δ and Bias are the merger price change and absolute value of the static prediction bias, respectively. Market share is the aggregate market share. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

To provide directional guidance on static model bias in mergers, we look at price changes as a function of pre-merger margins and shares, which are often used to simulate or approximate unilateral merger price increases (Miller et al., 2016). Columns (1)-(3) of Tables 13 and 14 explore how the percentage price change from a merger relates to pre-merger margins and market shares, which are often directly observed, as well as to primitives of the demand model. As is typically the case in static models, both pre-merger shares and margins are positively related to the size of the price change. Conditional on these observables, however, the dynamic parameters dampen the effect of a joint pricing merger but increase the effect of a brand consolidation merger. Correspondingly, the absolute value of the static model bias, which is the dependent variable in columns (4)-(6) of Tables 13 and 14, increases in joint pricing mergers but decreases in brand consolidation mergers as the pre-merger market share increases. Therefore, even if affiliation cannot be directly estimated, price change estimates should be revised accordingly if affiliation is expected to play an important role.

Table 14: Brand Consolidation Simulation: Merger Price Change and Bias

	(1) Price Δ	(2) Price Δ	(3) Price Δ	(4) Bias	(5) Bias	(6) Bias
Pre-Merger Market Share	21.874*** (0.218)	21.389*** (0.215)	28.305*** (1.337)	-9.580*** (0.174)	-10.170*** (0.166)	2.505*** (0.930)
Pre-Merger Margin	24.937*** (0.115)	24.856*** (0.112)	25.640*** (0.229)	5.476*** (0.092)	5.377*** (0.087)	6.917*** (0.193)
λ		1.360*** (0.076)	0.978*** (0.096)		1.656*** (0.059)	0.590*** (0.070)
α			-0.442*** (0.088)			-0.244*** (0.062)
ξ			-0.346*** (0.086)			-0.305*** (0.060)
$\bar{\xi}$			0.332*** (0.010)			0.236*** (0.008)
Price Δ						-0.235*** (0.008)
Constant	-3.919*** (0.046)	-4.323*** (0.050)	-7.044*** (0.243)	1.515*** (0.037)	1.023*** (0.039)	-0.639*** (0.168)
N	6566	6566	6566	6566	6566	6566

Notes: The dependent variables are outcomes from a merger simulation that consolidates products one and two under one brand. Observations are for firm 1. Price Δ and Bias are the merger price change and the absolute value of the static prediction bias, respectively. Market share is the aggregate market share. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B Additional Details on Merger Analysis

B.1 Implementation of Brand Consolidation Mergers

To implement the brand consolidation merger we adjust the ξ of the remaining product of the merging firm, such that it will retain the same share of unaffiliated customers if it were to keep the same pre-merger price. Then, given the adjustment, ξ_a , we allow the new single brand firm to price optimally. We therefore make the following adjustment:

$$\xi_a = \ln\left(\frac{s_1 + s_2}{s_0}\right) - \alpha p \quad (29)$$

Here, $s_1 + s_2$ is the combined pre-merger shares of products 1 and 2, s_0 is the share of the outside good, and p is the symmetric price of products 1 and 2.

B.2 Joint Pricing and Brand Consolidation Mergers in a Symmetric Static Logit Model

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 l into two firms yields equal shares for those two firms: $s_l(\bar{p}) = 2s_j(\bar{p})$

It follows that $s_l(p) = 2s_j(p) \quad \forall p$.

Proof: By construction:

$$s_l = \frac{e^{\delta_l - \alpha \bar{p}}}{1 + e^{\delta_l - \alpha \bar{p}} + \sum_g e^{\delta_g - \alpha p_g}} = 2 \frac{e^{\delta_j - \alpha \bar{p}}}{1 + 2e^{\delta_j - \alpha \bar{p}} + \sum_g e^{\delta_g - \alpha p_g}} = 2s_j \quad (30)$$

Equation (1) is true if and only if $e^{\delta_l - \alpha \bar{p}} = 2e^{\delta_j - \alpha \bar{p}}$. Dividing both sides by $e^{\delta_j - \alpha \bar{p}}$ and then taking logs, we have $\delta_l - \alpha \bar{p} - \delta_j + \alpha \bar{p} = \log(2)$. Therefore, $\delta_l - \delta_j = \log(2)$. It follows that if (i) and (ii) are true then $s_l(p) = 2s_j(p) \quad \forall p$.

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

$$\frac{d\Pi}{dp_1} = 1 + \alpha(p_1 - c_1)(s_1 - 1) + \alpha \sum_{k=2}^N (p_k - c_k)s_k = 0 \quad (31)$$

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 (2) simplifies to:

$$\frac{d\Pi}{dp_n} = 1 + \alpha(p - c)(s - 1) + \alpha(N - 1)(p - c)s = 0 \quad (32)$$

Solving equation (3) for p yields the symmetric, profit-maximizing price for each product:

$$p^* = \frac{1}{\alpha} \left[\frac{1}{1 - Ns} \right] + c \quad (33)$$

Now, suppose a single product firm sets a profit-maximizing price of p^* and therefore has a market share s_l . 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 $s_j(p^*) = \frac{s_l(p^*)}{2}$. By equation (4), the two product firm will set the same profit maximizing price, p^* . We therefore prove Lemma 2.

Lemma 1 and Lemma 2 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 Post-merger the firm retains both products and prices them to jointly maximize post-merger profits.
- ii Post-merger the firm removes one product from the market, say product 2. At the pre-merger price, the post-merger share of the first product would have the same market share as the sum of the two single product pre-merger firms. 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).

C Reduced-Form Evidence: Supplemental Results

C.1 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 Stock 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.⁴⁶ Using station-specific wholesale costs, we regress the 30-day lead wholesale cost on the current wholesale cost and the 30-day ahead future. In particular, we estimate the following equation.

$$c_{nt+30} = \alpha_1 c_{nt} + \alpha_2 F_t^{30} + \gamma_n + \epsilon_{nt} \quad (34)$$

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 overfit; 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 subsection 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.

⁴⁶Futures 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.

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 anytime 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 being reasonable in this context. We see no relationship between unexpected costs or expected costs and the price 30 days prior.⁴⁷

⁴⁷We interpret slight deviations from a zero as arising from an underlying correlation in unobserved cost shocks.

D Empirical Application: Supplemental Tables and Figures

D.1 Summary Statistics by Observation

Table 15: 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	0.0003	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	0.206	0.114	-0.440	0.132	0.272	1.048	110,844
Num. Stations	5.050	6.780	1	2	6	79	110,844

Notes: Table provides summary statistics for the observation-level data in the analysis. 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 US oil production data to generate an instrument for price in the demand estimation. Negative price-cost margins occur in 2.7 percent of observations.

D.2 Summary Statistics by Brand


Table 16: Summary of Brands

	Brand	Cond. Share	Share	Num. Markets	Num. Stations	Margins
1	Marathon	0.18	0.10	134	5.30	0.21
2	Sheetz	0.18	0.03	37	1.70	0.17
3	Speedway	0.17	0.03	39	3.70	0.18
4	Wawa	0.16	0.01	22	3.20	0.12
5	Exxon	0.15	0.07	116	4.60	0.25
6	7-Eleven	0.13	0.02	41	6.80	0.18
7	Shell	0.13	0.09	163	4.20	0.22
8	FRINGE	0.13	0.12	233	8.70	0.19
9	Pilot	0.12	0.01	21	1.40	0.13
10	BP	0.12	0.06	124	3.30	0.21
11	Loves	0.11	0.01	15	1.00	0.20
12	Valero	0.11	0.03	59	3.50	0.21
13	Thorntons	0.11	0.00	9	5.90	0.14
14	Hucks	0.11	0.00	10	1.90	0.15
15	Sunoco	0.09	0.01	32	4.70	0.28
16	Citgo	0.07	0.01	34	4.00	0.25

Notes: Table provides summary statistics by brand. The FRINGE brand is a synthetic brand created by aggregating brands that do not appear in 10 or more of the 241 markets in our data. Additionally, if a brand does not make up more than 2 percent of the average shares within a market, or 10 percent of the shares for the periods in which it is present, we also designate the brand as a fringe participant for that market.

D.3 Shares, Prices, and Instrument

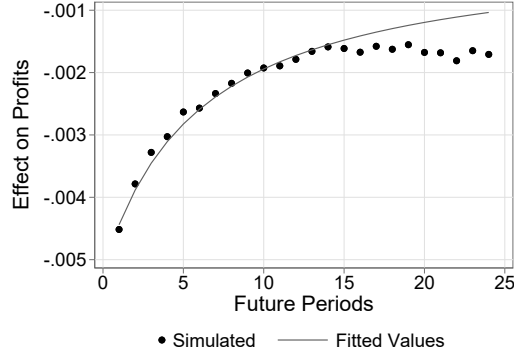
Figure 7: Shares and Prices



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.

E Details on Forward Simulations

Figure 8: 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 evaluate whether realized profits 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 specific week in a specific market by lowering that brand’s price by \$0.01. We re-compute shares in that period, then we compute equilibrium prices over the next 24 weeks. We repeat this experiment for all brands within the market and the period, and we iterate over markets and periods. We limit our analysis to the 75 markets and 52 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 $\beta = 0.998$, which corresponds to an annual discount rate of 0.901. We use the discounted values of $\frac{\partial \pi_j(t+\tau)}{\partial p_{jt}}$ to obtain an estimate of the full value of $\beta \frac{\partial E[V_j(\cdot)]}{\partial p_{jt}}$.

When calculating the net present value of the effect on profits, we use a fitted model to mitigate the effect of simulation noise. Specifically, we fit a curve of the form $\frac{\partial \pi_j(t+\tau)}{\partial p_{jt}} = \omega_1 \frac{1}{\tau + \omega_2}$. We set $\omega_2 = 6$ and we estimate $\omega_1 = -0.0311$ with least squares regression. In the regression, we include only up to 20 weeks into the future. Figure 8 plots the simulated effects on future profits against the fitted values. The figure shows that an increase in price today has a negative effect on future profits, and this effect shrinks over time. The fitted values explain 99 percent of the variation in mean values through $\tau = 20$. After week 15, the effect on future profits does not depreciate as much in the simulation as it does in our fitted model.