An Empirical Model of Dynamic Price Competition*

Alexander MacKay
Harvard University†

Marc Remer
Swarthmore College‡

September 29, 2017

**PRELIMINARY AND INCOMPLETE:
NOT FOR GENERAL DISTRIBUTION**

Abstract

Prices in many industries adjust slowly to changes in marginal costs. We consider the implications of price dynamics in the context of demand estimation, merger simulation, and cost pass-through. For the analysis, we develop an empirical model of dynamic consumer demand with oligopolistic competition. The model builds on the standard discrete-choice model and requires no supply-side assumptions to estimate, in contrast to existing methods, such as Bajari et al. (2007). We estimate the model using data on retail gasoline, where we find strong evidence of dynamic pricing. We demonstrate that welfare analysis depends critically on properly accounting for demand dynamics. Ignoring dynamic behavior results in upwardly biased predictions of merger price effects.

*An earlier version of this paper was entitled "The Dynamics of Pass-through with Expected and Unexpected Costs."

†Harvard University, Harvard Business School. Email: amackay@hbs.edu.

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

1 Introduction

Prices in a wide array of industries have been documented to react slowly to cost changes. Grocery stores (Peltzman, 2000), retail gasoline (Borenstein et al., 1997), and banking (Neumark and Sharpe, 1992) are just a few examples of markets where firms gradually pass through marginal cost shocks to consumers. While dynamically adjusting prices are an important feature of these markets, the demand estimation literature has largely ignored this property. The consequences of ignoring dynamics are three-fold. First, demand estimation may rely upon misspecified static optimization conditions, which can bias parameter estimates and fail to capture first-order incentives for policy changes. Second, static models will be unable to reproduce relevant short-run price dynamics, which can have important consumer welfare consequences in markets with frequently changing costs. Finally, pass-through estimates may fail to properly capture the completeness of the price adjustment when anticipation is not accounted for.

In this article, we develop and estimate an empirical model of dynamic demand. In response to demand-side behavior, profit-maximizing firms adjusting prices over many time periods in response to a cost change. The model also produces other important pricing patterns, such as price responses prior to anticipated future cost changes. The demand-side model includes a subset of consumers that become attached to the firm from which they purchased in the previous period, creating an incentive for firms to sacrifice current profits for future gains from the attached consumers. Our model of attachment can be used to describe brand loyalty, habit formation, lock-in, or search costs. This feature, which we call "lock-in" for the remainder of the paper, leads to firms in oligopolistic markets slowly adjusting prices to changes in cost.

The demand model is a straightforward extension of the standard discrete-choice logit model common in industrial organization. In our steady-state analysis, we show that optimal prices may either be higher or lower in the dynamic model compare to a static model, depending on the relative price sensitivity of locked-in consumers to consumers that are not locked-in. Given the parameters we estimate in the empirical application, we find that lock-in leads to higher markups in retail gasoline markets.

We consider mergers in the context of our model using numerical simulations in a duopoly setting. First, we demonstrate that the presence of lock-in in duopolistic competition may increase prices more than a merger to monopoly would in a static model. Second, we show that the relative price effect of a merger has non-monotonic patterns,[1] which is relevant for the consideration of antitrust authorities. The percent price effect of a merger may either increase or decrease with the lock-in probability. When consumers are sufficiently price insensitive, then lock-in allows firms to capture a large portion of the monopoly rents

[1]Absolute prices increase monotonically with the degree of lock-in.

pre-merger, resulting in a smaller impact of the merger on prices.

Next, we consider the empirical implications of failing to account for dynamic demand in a merger analysis. In our exercise, we calibrate a static demand model to data generated by our dynamic model and perform a merger simulation. Compared to the true impact of the merger, the (incorrect) static logit model systematically over-predicts merger price effects. In the dynamic model, the incentive to invest in future demand pushes prices down, and this effect remains after the merger. The static model falsely attributes a portion of this incentive to competition, which disappears post-merger, resulting in higher prices. Our analysis of mergers builds on the existing theoretical work on dynamic price competition (see, e.g., Farrell and Shapiro (1988), Beggs and Klemperer (1992), and Bergemann and Välimäki (2006)). Hendel and Nevo (2013) estimate an empirical model with dynamic pricing and discuss biases that could arise in a merger context. In relation to their paper, one of the distinguish factors is that we focus on settings with consumer loyalty, where purchases are positively correlated over time, rather than their setting with consumer stockpiling.

Our model can be taken to data, which we demonstrate through an empirical analysis of retail gasoline markets. We advance the empirical literature by developing an estimation method that accommodates dynamic demand, but does not require the assumptions used in standard dynamic estimation techniques, such as those developed by Bajari et al. (2007) and Pakes et al. (2007), to identify the dynamic parameters. Importantly, we do not require supply-side assumptions about the competitive game or the expectations of the firms to estimate demand. As supply-side behavior depends crucially on whether future shocks are expected or unexpected, we view this as a significant advantage.

The data needed to estimate the model are commonly used in the demand estimation literature, namely, prices, shares and cost shifters (or, alternatively, product characteristics). Nonetheless, we show that with firm-level shares (across all customer types), it is possible to separately identify each firms' current shares of locked-in and non-locked-in customers, as well as the probability that a consumer becomes locked-in after purchasing. Identification of locked-in customers allows for the presence of a serially correlated state variable for each firm. In other words, we estimate a model with endogenous unobserved heterogeneity. This flexibility has traditionally been a challenge for the estimation of dynamic models.

We estimate the model using a rich panel data set of prices, shares, and costs for retail gasoline stations. In this context, the model is best interpreted as one of habit formation or search costs, wherein some consumers return to the gas station from which they previously purchased without considering alternative sellers. We estimate that 76 percent of consumers only consider the gas station from which they previously purchased on a week-to-week basis, without searching competitors. On average, 98 percent of locked-in customers purchase gasoline, compared to only 34 percent of customers that are not locked-in ("free agents"). Locked-in customers display a much lower price sensitivity. Free-agent customers have an average elasticity of -8.1, whereas locked-in customers have an average elasticity of -0.03.

To highlight the importance of accounting for dynamic factors when estimating per-period firm payoffs, we impose a supply-side model, and we use the model-estimated static component of profits to estimate the dynamic component of the firms' first-order conditions. We project these estimates onto state variables to estimate how the dynamic incentives vary over time. Using this estimated function, we perform a merger analysis between two of the brands in our data. Relative to a static analysis, we find small effects from the merger. The merger increases both the static incentive to increase prices and the dynamic investment incentive to lower prices. Thus, the dynamic incentives serve to mitigate the static effects that arise from a merger.

Prior to developing our model, we present reduced-form evidence on dynamic pricing, in terms of cost pass-through, in retail gasoline markets. We use the data to separate out expected and unexpected costs, and we show that firms respond differentially to these two measures. Importantly, firms begin raising prices in anticipation of higher costs approximately 28 days prior to a cost shock. Firms exhibit "full" pass-through for expected cost changes, raising prices by 1.02 dollars for each dollar increase in cost. We find that the the pass-through of unexpected costs is limited, with firms raising prices by only 0.66 cents for an unexpected dollar increase in costs. We explore heterogeneity in firm pass-through and how it relates to the competitive conditions in the market.

2 A Oligopoly Model of Dynamic Demand

2.1 Demand Side

We develop a dynamic model of oligopolistic competition with product differentiation where consumers may become locked in to the firm from which they previously purchased. Lock-in may be interpreted as customer loyalty, habit formation, or a cost of search. Customers 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; rather, some consumers are likely to return to the same gas station due to habit formation or costly search behavior. As detailed below, we introduce differentiated product demand using the standard logit model, and then place the demand model into a dynamic oligopoly setting. The dynamics stem from firms current sales affecting future profits through the accumulation of locked-in consumers. We then demonstrate how common properties of cost pass-through (such as sticky prices, adjustment to expected cost shocks, and estimates of long-run pass-through) differ in this dynamic setting as compared to static environments.[2]

[2]Slade (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

The demand-side is characterized by the standard logit model,[3] but a fraction of consumers, δ, become locked-in to the firm from which they purchased in the previous period. We assume that consumers that purchase the outside good in period $t-1$ are never locked-in at time t. Consumers that are not locked-in, which we will refer to as "free agents," have utility $u_{jt}^{(n)} = \xi_j - \alpha p_{jt} + \eta_{jt} + \epsilon_{jt}^{(n)}$ from purchasing from any inside firm $j \in \{1, .., J\}$ or the outside good, $j = 0$. The superscript (n) indicates customer-specific values. For locked-in consumers, the utility from purchasing good j conditional on purchasing good i in the previous period is:

$$u_{jt}^{(n)} = \xi_j + \mathbf{1}[j = i]\bar{\xi} - (\alpha + \mathbf{1}[j = i]\bar{\alpha})p_{jt} + \eta_{jt} + \epsilon_{jt}^{(n)}. \tag{1}$$

Locked-in customers therefore experience a relative level utility shock $\bar{\xi}$ and a relative price sensitivity shock $\bar{\alpha}$. The relative utility shock, $\bar{\xi}$, may be interpreted as a switching cost, a brand-loyalty effect, or a cost of search where "non-searches" are more likely to return to their previously visited gas station without checking other prices.

For both free agents and locked-in consumers, we normalize the utility of the outside good to be zero. In our model, we include a product-time shock η_{jt}, which allows for potential correlation in unobservable utility shocks over time.

We denote the choice probability of a free-agent as $s_{jt}(0)$ and locked-in consumer as $s_{jt}(i)$, where i denotes the firm to which a customer is locked in and 0 denotes that a consumer is a free-agent. For ease of notation, we define $\xi_{jt} = \xi_j - \alpha p_{jt} + \eta_{jt}$ and $\xi_{jt} + \bar{\xi}_{jt} = \xi_j + \bar{\xi} - (\alpha + \bar{\alpha})p_{jt} + \eta_{jt}$. Then, given the standard logit assumption of a type 1 extreme value distribution on the utility shock, $\epsilon_{jt}^{(n)}$, the market shares of free agents and locked-in consumers are:

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

$$s_{jt}(i) = \frac{\exp(\xi_{jt} + \mathbf{1}[j = i]\bar{\xi}_{jt})}{1 + \sum_{k/i} \exp(\xi_{kt}) + \exp(\xi_{it} + \bar{\xi}_{jt})}. \tag{3}$$

The observable share for firm j is then given by the average of these two components, where each component is weighted by the relative proportion of customers in the free-agent or the locked-in state.

adjustments.

[3] In future versions, we plan to generalize the model to either the nested or random coefficients logit model.

2.2 Supply Side

Given the demand model above, firm j's profit-maximization problem can be written as follows:

$$\max_{p_{jt}}\{E[\sum_{t=0}^{\infty}\beta^t\Pi(p_t,r_t,c_t)]\}.$$

In this equation, p_t, r_t, and c_t are vectors of each firm's prices, total number of locked-in consumers, and marginal costs, respectively, at time t. Thus, firms compete by choosing prices that maximize expected discounted profits, where firms anticipate both future marginal costs and how current choices affect the future state of locked-in consumers.

The profit-maximization problem can be rewritten in Bellman equation form. To do so, we first define firm j's total share as:

$$S_{jt} = (1-\sum_k r_{kt})s_{jt}(0) + \sum_{i=1} r_{it}s_{jt}(i). \tag{4}$$

Thus, a firm's total share of sales can be written as a weighted sum of its share of free agent, $s_{jt}(0)$, and locked-in consumers, $s_{jt}(i)$. Note that firm j will make sales to consumers locked-in to other firms $j \neq i$, but the probability that such consumers will choose firm j is strictly lower than the choice probability of a free-agent consumer. We then define the number of consumers locked-in to firm j at $t+1$ to simply be $r_{jt+1} = \delta S_{jt}$. A fraction of a firms previous period consumers, δ, become locked-in next period, and therefore today's sales affect tomorrow's profits. This is captured in the following Bellman equation:

$$V_j(r,c) = \max_{p_{jt}|p_{kt\neq jt}} \big\{(p_{jt}-c_{jt})S_{jt} + \beta E(V_j(r_{t+1},c_{t+1})|p_t,r_t,c_t)\big\}. \tag{5}$$

In this Bellman formulation, firms optimal prices depend upon the number of consumers currently locked-in to each firm, r_{jt} and each firm's marginal cost c_{jt}, which enter as state variables. We therefore must solve for J optimal pricing policies conditional on $2J$ state variables. Note also that the optimal choice of prices depends essentially on the expectations of future costs, c_{t+1}. One contribution of the empirical work is to disentangle expected and unexpected costs to shed light on the sensitivity of dynamic welfare analysis to the proper accounting of expectations.

2.3 Steady-State Analysis

2.3.1 Monopoly

While solving for the optimal pricing policy in oligopoly markets is analytically intractable, we analyze steady-state prices in a monopoly market (with an outside good) to develop intuition about the impact locked-in consumers have on optimal prices and markups.

In the steady state, $r_{jt} = r_{jt+1} = r_j$ and $c_{jt} = c_{jt+1} = c_j$. To simplify notation in the monopoly case, let the firm's share of locked-in and free-agent consumers be s_j and s_0, respectively, and its number of locked-in consumers be r_t. It follows that the steady-state number of locked-in consumers, r^{ss} is:

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

The steady-state number of locked-in consumers is increasing in the probability of becoming locked in, δ, and the (absolute value) of the gap between the choice probabilities of locked-in customers and free agents, $s_j - s_0$. Using the steady-state value of locked-in consumers, we can solve for the steady-state pricing function.

The steady state period value is:

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

This equation represents the monopolists discounted profits, conditional on costs remaining at its current level. Thus, profits are increasing in both δ and the difference in choice probabilities of locked-in and free-agent consumers. These results are straight-forward: lock-in is profitable. Also, note that a model with no lock-in is embedded in this formulation ($\delta = 0$ 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 pricing function:

$$
p^{ss} = c^{ss} + \underbrace{\frac{-s_0\,(1-\delta s_j + \delta s_0)}{\frac{ds_0}{dp}\,(1-\delta s_j) + \frac{ds_j}{dp}\delta s_0}}_{\text{m = markup of price over marginal cost}}. \tag{6}
$$

The second term, m, on the right-hand side of equation (6) 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 lock-in. 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 consumer lock-in? In the simple case of no lock-in, when $\delta = 0$ and $s_j = s_0$, equation (6) reduces to the solution to the static model $p^{ss} = c^{ss} - \frac{s_0}{ds_0/dp}$. Denoting the markup term with lock-in as m_l and the markup term without lock-in as m_n, let us compare these two terms at the solution to the static model:

$$m_l = -\frac{s_0\left(1 - \delta s_j + \delta s_0\right)}{\frac{ds_0}{dp}\left(1 - \delta s_j\right) + \frac{ds_j}{dp}\delta s_0} \gtreqless -\frac{s_0}{ds_0/dp} = m_n.$$

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

$$m_l > m_n \iff -\frac{\partial s_0}{\partial p} > -\frac{\partial s_j}{\partial p}. \tag{7}$$

As the term m is declining in p (and the LHS of the pricing function is increasing in p), a higher value for m_l implies higher prices, and therefore higher markups. Thus, if locked-in consumers 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, arising from the fact that there is not a direct mapping between a comparison of elasticities and the above condition. In our model, free-agent consumers are more elastic than locked-in customers, i.e. $-\frac{\partial s_0}{\partial p}\frac{1}{s_0} > -\frac{\partial s_j}{\partial p}\frac{1}{s_j}$. However, locked-in customers also have higher shares, so markups may be higher or lower with consumer lock-in. Consider the logit model, where $-\frac{\partial s_0}{\partial p} = \alpha s_0(1 - s_0)$ and $-\frac{\partial s_j}{\partial p} = (\alpha + \bar{\alpha})s_j(1 - s_j)$. Though $\alpha(1 - s_0) > (\alpha + \bar{\alpha})(1 - s_j)$, the locked-in share s_j may be high enough relative to s_0 to flip the inequality.[4]

Thus, when locked-in consumer are less elastic, there may be an incentive to lower prices, relative to the static model. The intuition for this result is akin to those summarized in Farrell and Klemperer (2007); with dynamic demand and lock-in, firms face a trade-off between pricing aggressively today and "harvesting" locked-in consumers in future periods. In the steady state, our model shows that either effect may dominate. In the empirical section, we estimate $(\alpha + \bar{\alpha})$ to be -2.54 and the share of locked-in customers to be close to one. In this case, our data support markups being higher in in the presence of lock-in.

This finding has important implications for counterfactual exercises, such as merger simulation. Failing to account for customer lock-in will result in inferring markups from incorrect first-order conditions, which is a critical input to merger simulation. Furthermore, antitrust agencies often infer elasticities from markups calculated using accounting data (see, Miller et al. (2013)); our result demonstrates that this will lead to incorrect elasticities and merger price predictions when lock-in (or brand loyalty) is important.

[4] In the logit model, markups increase in the presence of lock-in when $s_0 > 0.5$.

2.3.2 Lock-In and Mergers

We now introduce competition into the model, and use numerical methods to analyze the steady-state of a duopoly market. In doing so, we demonstrate how the magnitude of the lock-in effect, δ, affects prices, markups, and profits. Furthermore, we analyze the impact of a merger to monopoly and find that the percentage change in price is not monotonic in δ and may even decrease with δ. We also analyze the consequences of ignoring the lock-in effect when calibrating demand and performing a merger simulation, and find that assuming the standard logit model when lock-in is present causes a systematic overestimation of the true price effect of the merger.

Each firm is specified to sell a single product and maximize the expected discounted value of profits. Therefore firm 1's Bellman equation is as follows:

$$V_1(r_1, r_2) = \max_{p_1|p_2} \{(p_1 - c_1)((1 - r_1 - r_2)s_1(0) + r_1 s_1(1) + r_2 s_1(2)) + \beta(V_1(r_1', r_2'))\}. \quad (8)$$

Here 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 locked-in shares for each firm, we focus on Markov perfect equilibrium.[5] Firm 1's profit-maximizing first-order condition is then:

$$\frac{d\pi_1}{dp_1} + \beta(\frac{dV_1'}{dr_1'}\frac{dr_1'}{dp_1} + \frac{dV_1'}{dr_2'}\frac{dr_2'}{dp_1}) = 0. \quad (9)$$

Firm 2's first-order condition is defined analogously. Next, we specify the derivatives of equation (8) with respect to r_1 and r_2 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 two conditions are:

$$\frac{dV_1}{dr_1} = \frac{\frac{d\pi_1}{dr_1} + \frac{d\pi_1}{dp_2}\frac{dp_2}{dr_1} + \beta\frac{dV_1}{dr_2}(\frac{dr_2'}{dr_1} + \frac{dr_2'}{dp_2}\frac{dp_2}{dr_1})}{1 - \beta(\frac{dr_1'}{dr_1} + \frac{dr_1'}{dp_2}\frac{dp_2}{dr_1})} \quad (10)$$

$$\frac{dV_1}{dr_2} = \frac{\frac{d\pi_1}{dr_2} + \frac{d\pi_1}{dp_2}\frac{dp_2}{dr_2} + \beta\frac{dV_1}{dr_1}(\frac{dr_1'}{dr_2} + \frac{dr_1'}{dp_2}\frac{dp_2}{dr_2})}{1 - \beta(\frac{dr_2'}{dr_2} + \frac{dr_2'}{dp_2}\frac{dp_2}{dr_2})}. \quad (11)$$

In the steady-state, $\frac{dV_1'}{dr_1'} = \frac{dV_1}{dr_1}$ and $\frac{dV_1'}{dr_2'} = \frac{dV_1}{dr_2}$. Therefore, we can plug equation (11) into equations (9) and (10) to eliminate $\frac{dV_1}{dr_2}$ and then equation (10) into equation (9) to eliminate $\frac{dV_1}{dr_1}$. This yields a steady-state profit-maximizing condition for firm 1 and the analogous one can be derived for firm 2. Then, in the steady-state , $r' = r$, and therefore

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

the equations that govern the evolution of locked-in consumers, $r'_j = \delta S_j$, can be leveraged to yield two more equilibrium restrictions:

$$r_1 = \frac{s_1(0) + r_2(s_1(2) - s_1(0))}{\frac{1}{\delta} + s_1(0) - s_1(1)} \tag{12}$$

$$r_2 = \frac{s_2(0) + r_1(s_2(1) - s_2(0))}{\frac{1}{\delta} + s_2(0) - s_2(2)}. \tag{13}$$

Equations (12) and (13) in conjunction with the steady-state profit-maximization conditions yield four restrictions that facilitate solving for the equilibrium prices and locked-in shares. However, there are four additional unknowns embedded in the envelope conditions: $\frac{dp_j}{dr_k}$ for $j, k \in \{1, 2\}$. These values are determined by the model, and we solve for them numerically using a local approximation method. For details, see the Appendix.

Finally, to produce results we must parameterize the demand model and specify values for marginal costs. The demand system is parameterized as follows:

Table 1: Duopoly Numerical Parameters

Variable	Name	Value
δ	Lock-in Probability	$\{0.05, 0.1, ..., 0.95\}$
α	Price Coefficient	$\{4, 4.1, ..., 7\}$
$\bar{\alpha}$	Price Lock-in Effect	-2.762
ξ_1, ξ_2	Intercepts	15.5, 15.7
$\bar{\xi}$	Intercept Lock-in Effect	2.218
c_1, c_2	Marginal Costs	2.2, 2.2
β	Discount Factor	0.96

Figure 1 plots the steady-state pre and post-merger prices for product 1.[6] To simulate the merger, we assume all demand parameters and marginal costs remain the same and only the ownership structure changes. Prices are depicted for values of δ ranging from 0.05 to 0.95 and $\alpha = 5.6$. We find that prices rise as the lock-in effect becomes stronger, both in the pre-merger duopoly and post-merger monopoly settings. Thus, in the steady-state, the "harvesting" effect dominates the incentive to invest in future demand. The pre-merger price for $\delta = 0.95$ is higher than the post-merger price for $\delta = 0.05$, indicating that lock-in may have a greater impact on price than reduced competition.

Table 2 summarizes the numerical results. Again, steady-state prices and margins appear to monotonically increase with the probability of consumers becoming locked-in. While this is an intuitive result in the case of a monopoly,[7] it also demonstrates that intense competition for future locked-in consumers does not dominate the incentive to increase prices to a loyal

[6]The prices for both products are approximately equal.

[7]However, Dube et al. (2009) demonstrates numerically that lock-in may lead to lower equilibrium prices, and indeed does so in its empirical application.

Table 2: Duopoly Lock-in Prices and Margins

Lock-in Probability (δ)	0.1	0.25	0.5	0.75	0.9
Pre-Merger Price					
$\alpha = 4$	2.70	2.75	2.89	3.17	3.46
$\alpha = 5$	2.59	2.62	2.71	2.84	2.98
$\alpha = 6$	2.47	2.49	2.53	2.59	2.67
$\alpha = 7$	2.37	2.37	2.39	2.42	2.48
Post-Merger Price					
$\alpha = 4$	3.68	3.72	3.80	3.95	4.16
$\alpha = 5$	3.03	3.05	3.12	3.23	3.39
$\alpha = 6$	2.63	2.65	2.69	2.78	2.90
$\alpha = 7$	2.40	2.41	2.44	2.49	2.58
$\%\Delta$ *Price from Merger*					
$\alpha = 4$	36.1	34.9	31.7	24.9	20.1
$\alpha = 5$	17.0	16.5	15.2	13.6	13.8
$\alpha = 6$	6.3	6.3	6.5	7.1	8.6
$\alpha = 7$	1.7	1.8	2.1	2.9	4.0
Pre-Merger Margin					
$\alpha = 4$	0.19	0.21	0.24	0.31	0.37
$\alpha = 5$	0.16	0.17	0.19	0.23	0.27
$\alpha = 6$	0.12	0.12	0.14	0.16	0.18
$\alpha = 7$	0.08	0.08	0.09	0.10	0.12
Post-Merger Margin					
$\alpha = 4$	0.41	0.41	0.43	0.45	0.48
$\alpha = 5$	0.28	0.29	0.30	0.33	0.36
$\alpha = 6$	0.17	0.18	0.19	0.21	0.25
$\alpha = 7$	0.09	0.10	0.11	0.13	0.16

Notes: Statistics are for product 1. α is the price coefficient for free-agent customers. The price coefficient for locked-in customers is $\alpha + \bar{\alpha}$.

Figure 1: Duopoly Steady-State Prices

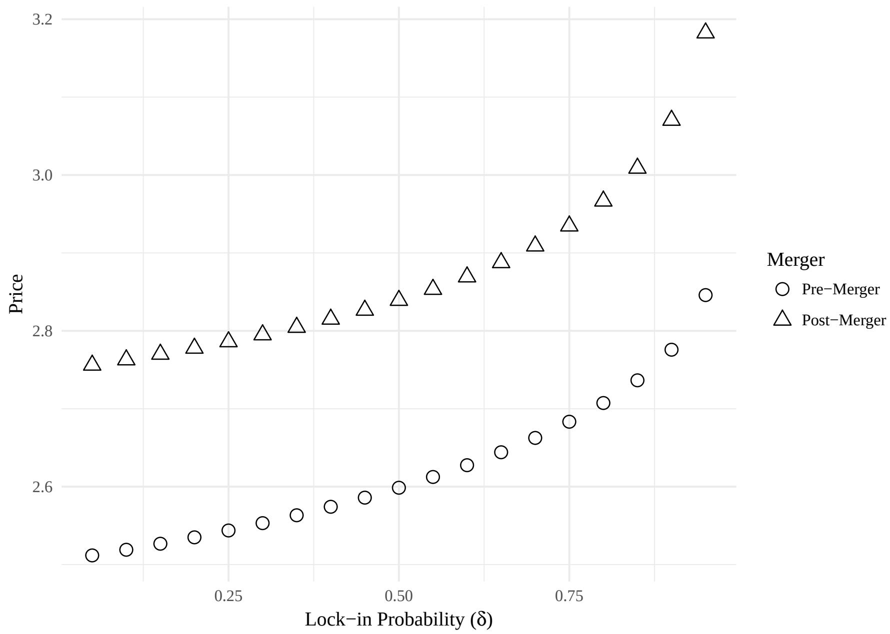

customer base. Also, prices and margins in Table 2 decrease as consumers become more price sensitive.

Interestingly, the the percentage price increase from a merger is not monotonic in the probability of consumers becoming locked-in. When consumers are less elastic (α = 4 or 5) the percentage price effects are generally large, ranging from 14 percent to 36 percent. At these values of α, however, the percent price change from the merger decreases with a stronger lock-in effect. On the other hand, when consumers are relatively more price sensitive, merger price effects increase with the lock-in effect. The reason for these results is, in part, due to the relative value of competing for versus harvesting locked-in consumers. When α is lower, holding $\bar{\alpha}$ constant, then locked-in consumers are less likely to choose the competing firm. Consequently, harvesting locked-in consumers is more valuable for a duopolist at lower levels of α. In turn, pre-merger prices and competition is less intense, and post-merger monopolists have less consumer rent to extract (on a percentage basis). On the other hand, when α is high, then pre-merger competition for free-agents is more intense and harvesting is relatively less valuable. Thus, increasing δ leads to relatively modest increases in pre-merger prices, and therefore the merger leads to greater price increases (on a percentage basis).

2.3.3 Model Misspecification in Merger Simulation

We now explore the implications of failing to account for consumer lock-in when calibrating demand and simulating a merger, as is common practice by antitrust practitioners. To do so, we consider the following hypothetical scenario. The true underlying model is the duopoly lock-in model. A practitioner observes both firms' pre-merger prices, marginal costs, and aggregate market shares (rather than separately observing its free-agent and locked-in shares). This data is then used to recover the demand parameters of the standard logit model (ξ_i and α), and then the price effects of a merger are simulated.[8]

Figures 2 (a) and (b) depict percent price increases from the true lock-in model (for product 1), and the predicted price increases from the standard logit model calibrated to match the observed pre-merger prices, marginal costs, and aggregate firm market shares. In these figures, α is set to be 4.8 and 6.4, respectively. In both figures, the logit model over-predicts the price increase and the bias worsens as the the lock-in effect increases. Thus, ignoring consumer lock-in becomes increasingly problematic as the effect strengthens.

These figures illustrate potential problems in antitrust enforcement if consumer lock-in is not accounted for in merger simulation. Antitrust agencies typically evaluate mergers based upon whether or not a transaction will result in a "small but significant non-transitory increase in price," which is often taken to be 5 percent.[9] In Figure 2 (b), for values of δ between 0.4 and 0.65, the lock-in model predicts prices increases below a 5 percent SSNIP, whereas the incorrectly specified logit model finds a price change above the threshold. In Figure 2 (a), the price predictions of the two models diverge and become increasingly disparate as δ increases. While the price increases are above a SSNIP in both models, different demand parameters could result in similar price paths centered at a lower price increase, say 5%. Thus, depending on the underlying demand parameters, it could be that the lock-in and standard logit models lead to different enforcement decisions across nearly the entire range of δ.

Table 3 provides simulation results for a broader range of parameters and confirms the patterns depicted in Figures 2 (a) and (b). The logit model calibrated to match pre-merger lock-in observables always over-predicts the true price effects. The magnitude of the bias (in terms of percentage points) increases with the lock-in effect and decreases with the free-agent price coefficient (α). To calibrate the logit model, the price coefficient is inferred from observed margins. The table demonstrates the logit coefficient is almost always calibrated to be more elastic than the share-weighted lock-in price coefficient (columns 4 and 3, respectively). Also, while not depicted in the table, the mean utility in the logit model is also calibrated to be lower when compared to the share-weighted counterparts in the lock-in

[8]See Miller et al. (2016) for details on the calibration and simulation procedure for the logit model.

[9]See the 2010 Horizontal Merger Guidelines issued jointly by the Federal Trade Commission and the US Department of Justice.

Figure 2: Simulated Merger Price Increases

(a) $\alpha = 4.8$

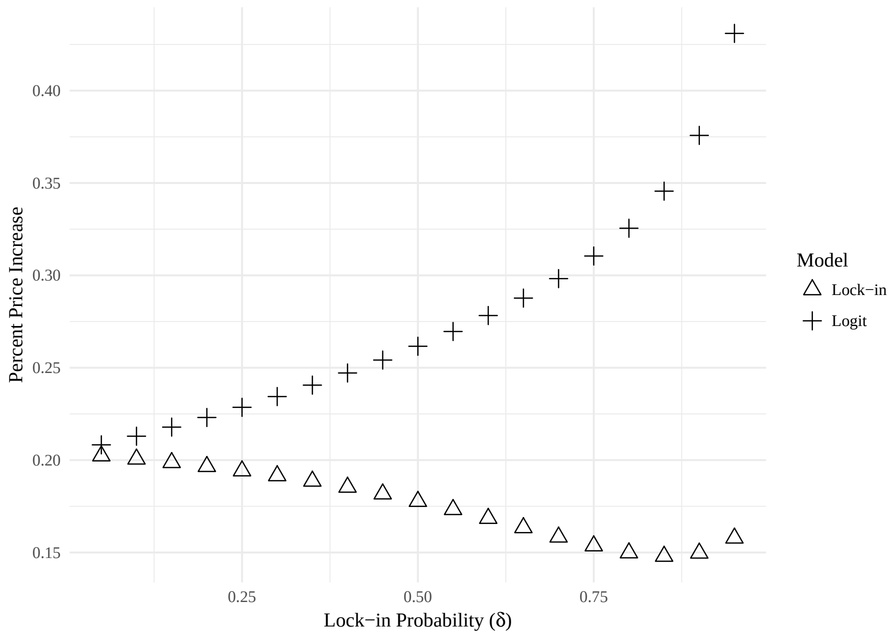

(b) $\alpha = 6.4$

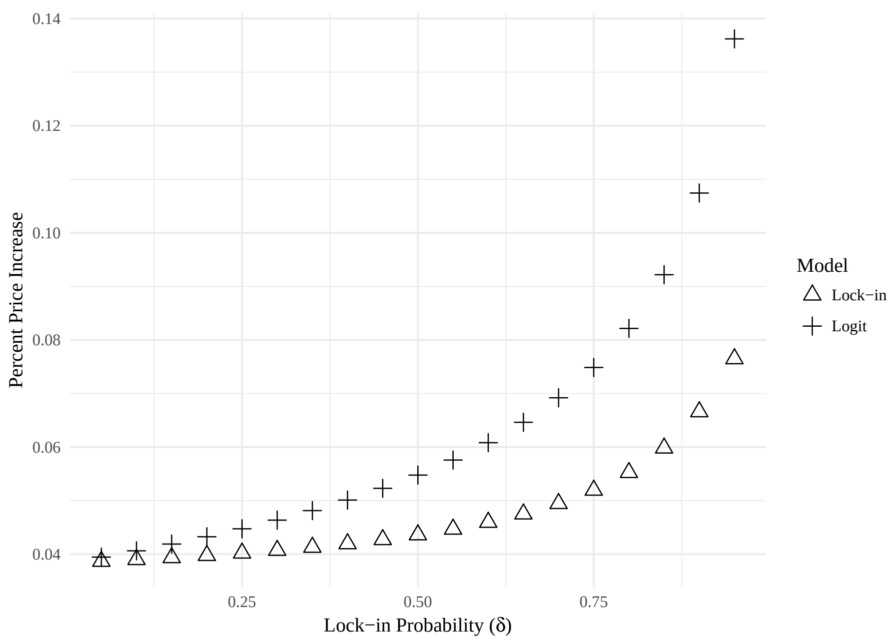

Table 3: Duopoly Merger Simulation: Lock-in vs. Logit

δ	Free-Agent α	Lock-in Avg. α	Logit α	Pre p_1	Lock-in Δp_1	Logit Δp_1	Bias
0.10	4	3.72	3.79	2.70	36.1	38.3	2.2
0.25	4	3.31	3.45	2.75	35.0	41.0	6.1
0.50	4	2.62	2.80	2.89	31.7	46.9	15.2
0.75	4	1.93	2.01	3.17	24.9	54.2	29.3
0.90	4	1.51	1.54	3.46	20.1	162.6	142.6
0.10	5	4.73	4.76	2.59	17.0	18.0	1.0
0.25	5	4.31	4.37	2.62	16.5	19.3	2.9
0.50	5	3.62	3.67	2.71	15.2	22.2	7.0
0.75	5	2.93	2.91	2.84	13.6	26.7	13.1
0.90	5	2.51	2.44	2.98	13.8	33.0	19.1
0.10	6	5.72	5.82	2.47	6.3	6.6	0.3
0.25	6	5.31	5.53	2.49	6.3	7.7	0.9
0.50	6	4.62	5.00	2.53	6.5	8.6	2.1
0.75	6	3.93	4.37	2.59	7.1	11.3	4.2
0.90	6	3.51	3.81	2.67	8.6	15.5	6.9
0.10	7	6.72	6.94	2.37	1.7	1.7	0.0
0.25	7	6.31	6.83	2.37	1.8	1.9	0.1
0.50	7	5.62	6.57	2.39	2.1	2.5	0.3
0.75	7	4.93	6.09	2.42	2.9	3.7	0.8
0.90	7	4.52	5.35	2.48	4.0	5.7	1.8

Notes: Statistics are for product 1. α is the price coefficient for free-agent customers in the lock-in model. The price coefficient for locked-in customers is $\alpha + \bar{\alpha}$, where $\bar{\alpha} = -2.762$. Lock-in Avg. α is the average price coefficient in the market, where the average is weighted by the shares of free-agent and locked-in consumers. Logit α is the calibrated price coefficient for the standard logit model. Bias is post-merger price prediction error measured in percentage points.

model. This occurs because, conditional on the observed margins, the logit model interprets the incentive to invest in future demand, which puts downward pressure on margins, as more elastic demand.

Yet, prices are always predicted to rise by more in the logit model. Biased predictions of the logit model are likely to arise from ignoring the dynamic incentive to invest in future demand, rather than a biased elasticity or mean utility parameters alone. In the static logit model, there is no incentive for the post-merger monopolist to invest in future demand, and therefore it sets its post-merger price simply to balance current marginal revenue and marginal cost. On the other hand, in the dynamic lock-in model, there is still an incentive for the monopolist to invest in future demand, which imposes downward pressure on post-merger prices. Thus, starting from the same pre-merger prices the logit model leads to greater post-merger price effects, even when its demand parameters are biased toward greater elasticity and lower mean utility.

3 Reduced-Form Evidence of Dynamics

To motivate the structural model, we first provide evidence on dynamically adjusting retail gasoline prices and dynamic demand. A host of previous studies have found that retail gasoline prices may take multiple weeks to fully incorporate a change in marginal cost.[10] 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. We therefore almost perfectly observe each gas station's marginal cost changes, except for privately negotiated discounts per-gallon, which are likely fixed over the course of a year.[11] The data ranges from September 25th, 2013 through September 30th, 2015. The data was obtained directly from the Oil Price Information Service (OPIS), which routinely supplies data used in academic studies (e.g. Lewis and Noel 2011; Chandra and Tappata 2011; Remer 2015).

OPIS also supplied the market share data, which it obtains directly from "actual purchases that fleet drivers charge to their Wright Express Universal card." The data is specified at the weekly, county/gasoline-brand level.[12] 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).

3.2 Identifying 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.[13]

[10] See Eckert (2013) for a comprehensive review of the literature.

[11] The data also include all federal, state, and local taxes.

[12] In 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.

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

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.

$$c_{it+30} = \alpha_1 c_{it} + \alpha_2 F_t^{30} + \gamma_i + \epsilon_{it} \tag{14}$$

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

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.

3.2.1 Note on the 30-Day Ahead Expectation

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. We do find a small anticipatory effect of our measure of unexpected costs on price 14 days in advance, but is less than 5 percent of the total price adjustment. We interpret this correlation to arise from the underlying correlation in unobserved cost shocks.

3.3 Dynamic Pricing: Pass-through Regressions

Further highlighting the temporal component of cost pass-though, 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 employ our measures of expected and unexpected costs to study how firms differentially respond to these costs. 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:

$$p_{it} = \sum_{s=-50}^{50} \beta_s \hat{c}_{it-s} + \sum_{s=-50}^{50} \gamma_s \tilde{c}_{it-s} + \sum_{s=-50}^{50} \phi_s \tau_{it-s} + \psi_i + \varepsilon_{it}. \tag{15}$$

Here, p_{it}, is the price observed at gas station i at time t. $\hat{c}_{it-s}$ and $\tilde{c}_{it-s}$ are the expected and unexpected wholesale costs observed with lag s, and τ_{it-s} is the state-level sales tax.[14] 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.[15]

Figure 3 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, reaching a peak of 0.71 after 34 days.

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 is reached 21 days after the shock. Though the total duration of adjustment is longer compared to the unexpected cost shock, the firm incorporates the cost more quickly after it is realized. This coincides with substantial anticipation by the firm; the price already captures about a third of the effect of the expected cost shock the day before it arrives.

A 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 66 cents for each dollar increase in cost.

The different response to unexpected and expected costs emphasizes the need for empirical researchers to think carefully about designing proper estimators for pass-through.

[14] To more easily incorporate future anticipated costs into the regression, we do not estimate an error-correction model (Engle and Granger, 1987), which is commonly used to estimate pass-through in the retail gasoline literature.

[15] As a robustness check, we also estimated the price response to expected and unexpected costs using the error-correction model, and we found nearly identical results.

Figure 3: Cumulative Pass-through

(a) Unexpected Costs

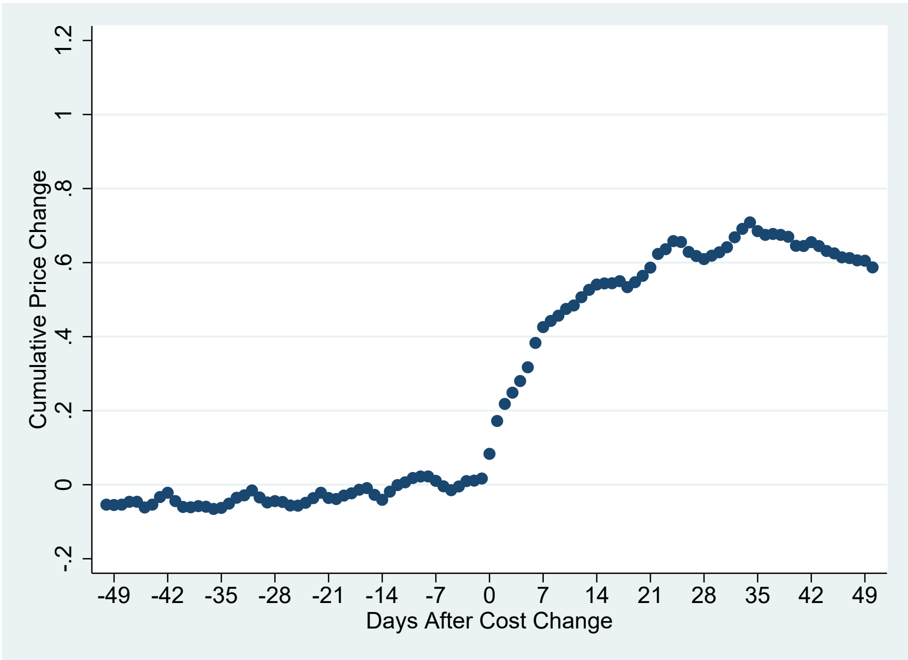

(b) Expected Costs

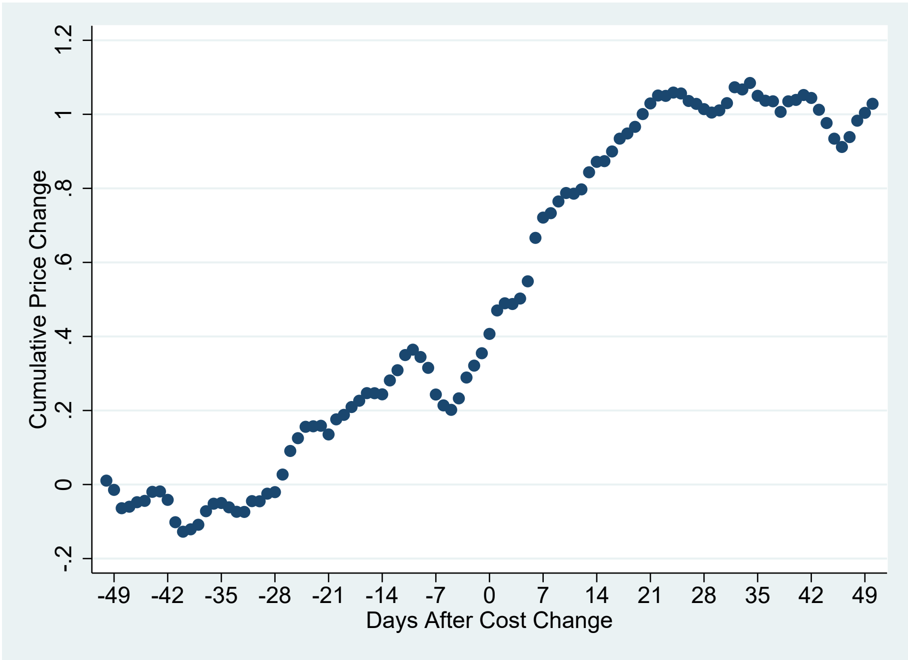

If costs are anticipated, then a pass-through measure that omits leads will only capture a portion of the overall response. Additionally, if a pass-through estimate is to be used for evaluation, is it important to ensure that the estimator relies on a mix of unexpected and expected shocks that translate to the policy under analysis. In our setting, for example, an analysis that used an unexpected cost shock to predict the impact of a tax rate change would be inappropriate, as the tax change is anticipated and leads to a much greater price response.

An important consideration for pass-through when analyzing imperfect competition is the distinction between idiosyncratic costs and common costs. In the current section, we consider only the simple cut between unexpected and expected costs to focus attention on this previously unexplored dimension of pass-through. In Appendix B, we present results for common costs and idiosyncratic costs. In our setting, costs are highly correlated, with common costs tending to dominate idiosyncratic costs at moderate frequencies. Therefore, the results for common costs are very similar to those in this section. One distinction, however, is that pass-through of unexpected common costs is higher than the pass-through for total unexpected costs. For idiosyncratic costs, which we consider by controlling for the prices of rivals, pass-through is on the order of 0.04 to 0.06. This is not surprising for a highly competitive market such as retail gasoline.

3.4 Heterogeneity

We now show that there is significant heterogeneity across firms in the dynamics of cost pass-through, and that it can, in part, be explained by the firm's competitive environment. To do so, we estimate station-specific regressions of equation (15), and then construct a cumulative response function for expected and unexpected costs separately for each gas station. To summarize these response functions, we use non-linear least squares to fit a three-segment spline to each firm's cumulative response function. We restrict the slopes of the first and last segments to equal zero. The middle segment captures the duration and magnitude of the price adjustment. This methodology allows to to estimate for each gas station (i) anticipation: how far in advance the firm begins responding to a cost shock, (ii) duration: how many days it takes to reach the new long-run equilibrium, and (iii) rate: what proportion of the cost change is passed through to price each day. We also construct the long-run pass-through rate ("LR_PTR"), which reflects the degree to which prices reflect costs at the end of the adjustment period.

Table 4 summarizes the estimated heterogeneity parameters. The five rows correspond to the response to unexpected costs, and the last five rows correspond to expected costs. These summary statistics align with the mean pass-through parameters estimated in the previous section. The median firm anticipate expected costs much earlier (25.4 days in advance) than unexpected costs (4.1 days in advance), and full pass-through is greater

Table 4: Pass-through Heterogeneity

Parameter	p10	p25	p50	p75	p90
Anticipation	-24.5	-15.0	-4.1	-0.1	3.6
Rate	0.009	0.016	0.026	0.045	0.103
Duration	5.0	12.3	26.9	47.0	59.4
Finish	4.0	10.7	21.5	32.6	41.6
LR_PTR	0.278	0.458	0.657	0.862	1.146
Anticipation_Exp	-39.0	-33.0	-25.4	-14.0	-4.0
Rate_Exp	0.015	0.017	0.021	0.029	0.051
Duration_Exp	19.2	33.2	48.5	60.0	70.7
Finish_Exp	10.2	15.8	23.0	30.0	38.0
LR_PTR_Exp	0.860	0.940	1.015	1.100	1.209

for expected costs (1.02) than unexpected costs (0.66). The median firm takes longer to incorporate expected costs, but tends to reach full pass-through around the same time for both types of shocks (21.5 versus 23.0 days after the shock).

Comparing the 10-90 percentile range of the estimates, we find that there is greater variation in the pricing response to unexpected costs. The median rate of adjustment is faster for unexpected costs (0.026/day compared to 0.012), but sees a spread of (0.009, 0.103) for the 10-90 percentile range, compared to (0.015, 0.051). Likewise, there is greater variation in the timing of when firms fully incorporate cost shocks in price and the full pass-through rates (0.278 to 1.146 for unexpected, 0.860 to 1.209). The long-run pass-through estimates for expected costs display a tendency toward homogeneity, as they are clustered around 1.

For the specification in this section, our fitted splines consist of three parameters, where we restrict the level of the first segment to be zero. In Appendix C, we consider another version that adds the level of the first segment as a fourth parameter. The results are consistent.

3.5 Competition and Heterogeneity Regressions

In this section, we project our estimated parameters onto a simple measure of market competition. In doing so, we demonstrate that competition has an important effect on the dynamics of cost pass-through, which further motivates the dynamic oligopoly model and demonstrates that firm expectation play an important role in determining cost pass-through rates. We perform a series of county-level regressions that relate the estimated parameters in Table 4 to the Herfindahl-Hirschman index ("HHI") in the county.

To do so, we first calculate the median firm-level parameter in each county, and regress it on the HHI and the number of gas stations in the the county.[16] To calculate the HHI, we sum the square of each gasoline brand's weekly county market share. We then take the average across all weeks for each county to use as an independent variable in the regression.[17] We also include the number of gas stations in each county as a regressor. This accounts for variation in population and demand across counties in a simple way; without this control, counties with a large consumer base would, all else equal, have lower HHI.

Table 5: Pass-Through and County Competition

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Variable	LR_PTR	LR_PTR_exp	Rate	Rate_exp	Anticipation	Anticipation_exp	Duration	Duration_exp
HHI	0.182	0.041	0.016	0.014**	-2.592	14.058**	-2.159	-16.221*
	(0.138)	(0.055)	(0.023)	(0.007)	(4.051)	(5.662)	(7.212)	(8.312)
Number of Stations	0.000	-0.000	0.000	0.000	-0.013	0.036	-0.002	-0.051
	(0.001)	(0.000)	(0.000)	(0.000)	(0.017)	(0.023)	(0.046)	(0.034)
Constant	0.629***	1.014***	0.034***	0.021***	-5.493***	-28.502***	29.190***	53.559***
	(0.044)	(0.019)	(0.008)	(0.003)	(1.415)	(1.848)	(2.744)	(2.758)
Observations	248	248	248	248	248	248	248	248

Notes: LR_PTR is the long-run full pass-through rate. "_exp" represents the reaction to an expected cost change, and other variables are reactions to unexpected cost changes. HHI is Herfindahl-Hirschman index. Obsevations are at the county-level. The dependent variable is the median estimated firm-level parameter in each county. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

The results of the series of regression are presented in Table 5, and demonstrate that competition affects pass-through dynamics. We find a significant relationship between the HHI and reaction to expected changes in cost. Specifically, the anticipated response ("Anticipation_exp"), the rate of adjustment ("Rate_exp"), and duration of adjustment ("Duration_exp") are all significantly related to the county HHI such that less competition leads to later and quicker responses to anticipated cost changes. Interestingly, we find in this simple regression that the long-run pass through rate ("LR_PTR") is unaffected by competition. Adding additional controls, however, such as population or mean income leads to positive and significant results, in some specifications. Thus, there is some evidence that less competition also leads to a greater proportion of expected cost changes to be passed-through to consumers On the other hand, we do not estimate a significant relationship between pass-through of unexpected cost shocks and the HHI. This underscores the importance of distinguishing between expected and unexpected cost changes when analyzing the interaction of competition and pass-through dynamics. In the following sections, we develop a structural model to capture these reduced-form pricing dynamics

[16]Results are qualitatively the same when using the mean parameter value. We choose the median, as the distribution of parameter estimates are slightly skewed and there are a few outliers.

[17]We take the weekly average HHI, as we observe market share rather than quantity.

Table 6: Regressions with Share as the Dependent Variable

	(1)	(2)	(3)	(4)	(5)
Price	0.011***	0.000**	0.004	−0.002	−0.072***
	(0.001)	(0.000)	(0.002)	(0.004)	(0.017)
Lagged Share		0.973***	0.963***	0.554***	0.628***
		(0.001)	(0.001)	(0.002)	(0.003)
Price Squared			−0.000	0.001	0.010***
			(0.000)	(0.001)	(0.003)
Comp. Price (Mean)			−0.004***	−0.006***	−0.106**
			(0.001)	(0.001)	(0.044)
Comp. Price (SD)			−0.001	−0.002	0.088***
			(0.001)	(0.001)	(0.024)
Comp. Stations			−0.000***	−0.000***	−0.004***
			(0.000)	(0.000)	(0.000)
Num. Stations			0.000***	0.004***	0.000
			(0.000)	(0.000)	(.)
Num. Brands			−0.001***	−0.002***	0.000
			(0.000)	(0.000)	(.)
Week FEs				X	
County-Brand FEs				X	
Brand-State-Week FEs					X
Week-County FEs					X
County-Brand-WofY FEs					X
Observations	175565	170935	170775	170756	156212
R^2	0.00	0.95	0.95	0.96	0.98

Standard errors in parentheses

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

3.6 Dynamic Demand: Correlation in Shares Over Time

Though ultimately the importance of demand-side dynamics in our data will be estimated by the model, it is informative to examine the reduced-form relationships between key elements. The dynamic model we develop in the next 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 6.

The regressions indicate that lagged shares are a significant predictor of current shares. 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, we include rich multi-level fixed effects: by 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.[18]

4 Empirical Application: Demand Estimation

We now present the empirical application of our 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 our methodology, we present results for demand estimation. In Section 5, we use our estimated demand system to analyze the dynamic incentives faced by suppliers. We use these results to consider a merger between two brands.

4.1 Demand Estimation Methodology

Given our dynamic extension of the logit demand system, we obtain the familiar expression for the log ratio of shares from Equation (2):

$$\ln s_{jt}(0) - \ln s_{0t}(0) = \xi_{jt} \tag{16}$$

Likewise, we obtain the following relation for the shares of locked-in customers:

$$\ln s_{jt}(i) - \ln s_{0t}(i) = \xi_{jt} + \mathbf{1}[j = i]\bar{\xi}_{it} \tag{17}$$

We can combine Equations (16) and (17) to obtain the following:

$$\ln s_{jt}(i) - \ln s_{0t}(i) - (\ln s_{jt}(0) - \ln s_{0t}(0)) = \mathbf{1}[j = i]\bar{\xi}_{it} \tag{18}$$

When we specify $\bar{\xi}_{it} = \bar{\xi} - \bar{\alpha}p_{it}$, then this empirical relationship depends only on the dynamic parameters $\bar{\xi}$ and $\bar{\alpha}$ along with the observed prices.

[18] We have also estimated specifications that add lagged prices. Though the first lag is significant, there is almost no effect on the lagged share coefficient.

4.1.1 Separating Locked-In Shares from Observed Shares

The challenge in dynamic demand estimation is that we do not separately observe the non-locked-in and free-agent customers. Instead, we observe the aggregate share, S_{jt}, which is a weighted combination of the $\{s_{jt}(i)\}$ and depends on the number of locked-in customers for each product $\{r_{jt}\}$. Observed shares are determined by the following:

$$S_{jt} = (1 - \sum_k r_{kt}) \cdot s_{jt}(0) + \sum_i r_{it} \cdot s_{jt}(i).$$

To separate out $s_{jt}(i)$ from S_{jt}, we leverage the structure of our model. First, we note that r_{jt} is proportional to the observed shares in the previous period: $r_{jt} = \delta S_{jt-1}$. Second, we use the relations in the previous section to obtain the following expressions:

$$s_{jt}(0) = \left(\frac{s_{0t}(0)}{s_{0t}(j)} - 1\right) \frac{1}{\exp(\bar{\xi}_{jt}) - 1}$$

$$s_{jt}(i) = s_{0t}(i) \frac{s_{jt}(0)}{s_{0t}(0)} \cdot \exp\left(\mathbf{1}[j = i]\bar{\xi}_{it}\right)$$

That is, the $J + J^2$ unknowns $\{s_{jt}(i)\}$, can be expressed in terms of the $J + 1$ unknowns $\{s_{0t}(j)\}$ and $s_{0t}(0)$.

The observed share equation gives us J restrictions:

$$\begin{aligned} S_{jt} = (1 - \sum_k r_{kt}) \cdot \left(\frac{s_{0t}(0)}{s_{0t}(j)} - 1\right) \frac{1}{\exp(\bar{\xi}_{jt}) - 1} \\ + \frac{s_{jt}(0)}{s_{0t}(0)} \cdot \sum_i r_{it} \cdot s_{0t}(i) \exp\left(\mathbf{1}[j = i]\bar{\xi}_{it}\right) \end{aligned}$$

And the final restriction, $1 - \sum_k s_{jt}(0) - s_{0t}(0) = 0$ identifies the shares $\{s_{jt}(i)\}$, conditional on the dynamic parameters.

4.1.2 Computational Simplicity

We can further reduce the burden of solving for the $J+1$ parameters $\{s_{0t}(j)\}$in each market by transforming the restrictions into a quadratic function of two parameters in each market.

Using the fact that

$$\sum_k r_{kt} \cdot s_{0t}(k) \exp\left(\mathbf{1}[j = k]\bar{\xi}_{kt}\right) = \sum_k r_{kt} s_{0t}(k) + \left(\exp(\bar{\xi}_{jt}) - 1\right) r_{jt} \boldsymbol{s_{0t}(j)}$$

We can write

$$
\begin{aligned}
0 = &\left[\exp(\bar{\xi}_{jt}) - 1\right] \frac{r_{jt}}{s_{0t}(0)} \boldsymbol{s_{0t}(j)}^2 \\
&+ \boldsymbol{s_{0t}(j)} \left(\left[\exp(\bar{\xi}_{jt}) - 1\right] (S_{jt} - r_{jt}) + \frac{1}{s_{0t}(0)} \sum_{0,k} r_{kt} s_{0t}(k) \right) \\
&- \sum_{0,k} r_{kt} s_{0t}(k)
\end{aligned}
$$

and solve for $\{s_{0t}(t)\}$ as a quadratic function of observable variables and the two parameters $s_{0t}(0)$ and $\sum_{0,k} r_{kt} s_{0t}(k)$.

4.1.3 Identification of the Dynamic Parameters

As we have shown above how to decompose shares for any values of the dynamic parameters, we need additional restrictions to identify the dynamic parameters. In order to do so, we assume that any idiosyncratic product-time shocks, after allowing for a period-specific fixed effect, are uncorrelated over time. These shocks may contain real demand shocks or measurement error from our data. That is, we assume $\eta_{jt} = \phi_t + \kappa_{jt}$, and $Cov(\kappa_{jt}, \kappa_{j(t+1)}) = 0$.

4.1.4 Implementation

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

1. First, pick values for δ, $\bar{\xi}$, and $\bar{\alpha}$.
2. Calculate $r_{jt} = \delta S_{jt-1}$ for all periods except the first.
3. In each market, solve for $s_{0t}(0)$ and $\sum_{0,k} r_{kt} s_{0t}(k)$ using the non-linear system of equations obtained previously. Find $s_{jt}(0)$ for each firm.
4. Run the regression implied by equation (16) using the $\{s_{jt}(0)\}$ obtained in the previous step. Calculate the correlation of the residuals $Cor(\hat{\kappa}_{jt}, \hat{\kappa}_{j(t+1)})$.
5. Repeat 1-5 to find the dynamic parameters that set the correlation to zero.

The regression for Equation (16) 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 solver only has to find two parameters, $s_{0t}(0)$ and $\sum_{0,k} r_{kt} s_{0t}(k)$, for each market-period. The quadratic form 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.

4.2 Data for Structural Model

We supplement our 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 four weeks. For any gap greater than four weeks, we assume the station was not in the choice set for that gap. 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.

To reduce the sensitivity of our analysis to brands with small shares and to make the counterfactual exercises more computationally tractable, we aggregate brands with small shares in our data into a synthetic "fringe" brand. We designate a brand as part of the fringe if it does not appear in ten or more of our 252 counties. Additionally, if a brand does not make up more than 2 percent of the average shares within a county, 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 county. These steps reduce the number of observations from 194,275 down to 112,417. Additionally, this reduces the maximum number of brands we observe in a county to 8, down from 24.

Table 7: Summary of Brands

	Brand	Cond. Share	Share	Num. Markets	Num. Stations	Margins
1	Marathon	0.18	0.10	135	5.20	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	0.16	0.07	119	4.60	0.25
6	Hucks	0.15	0.01	11	1.80	0.15
7	7-Eleven	0.15	0.02	42	6.70	0.18
8	FRINGE	0.14	0.12	244	9.90	0.19
9	Shell	0.13	0.09	165	4.20	0.22
10	Pilot	0.12	0.01	21	1.40	0.13
11	BP	0.12	0.06	127	3.30	0.21
12	Loves	0.11	0.01	15	1.00	0.19
13	Valero	0.11	0.02	59	3.50	0.21
14	Thorntons	0.11	0.00	9	5.90	0.14
15	Sunoco	0.10	0.01	35	4.30	0.29
16	Citgo	0.08	0.01	35	3.90	0.24

The resulting brands are displayed in Table 7. We reduce the number of brands for the analysis to 16. The FRINGE brand is, on average, 14 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).

Table 8: Summary Statistics by County

Statistic	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max	N
Num. Brands	4.42	1.51	1	3	5	8	252
Price	2.87	0.11	2.69	2.77	2.95	3.14	252
Wholesale Price	2.25	0.05	2.12	2.21	2.28	2.44	252
Margin	0.21	0.06	0.06	0.17	0.24	0.44	252
Num. Stations	22.65	28.52	1.12	7.00	25.48	243.47	252

Table 8 provides summary statistics of the data for our 252 counties. There is variation in the number of brands we observe in each county, ranging from 1 to 8. There is cross-sectional variation in wholesale prices, margins, and the number of stations in each county.

Table 9: Retail Gasoline in Kentucky and Virginia: Oct 2013 - Sep 2015

Statistic	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max	N
Share	0.141	0.110	0.0003	0.060	0.187	0.688	112,930
Price	2.870	0.529	1.715	2.383	3.310	4.085	112,930
Wholesale Price	2.256	0.527	1.245	1.754	2.673	3.545	112,930
Wholesale FE	2.260	0.031	2.205	2.229	2.291	2.365	112,930
Margin	0.205	0.115	−0.440	0.131	0.273	1.048	112,930
Num. Stations	5.256	7.165	1	2	6	83	112,930
Food	1.543	0.817	0.000	1.000	2.000	4.000	112,633
Supermarket	0.019	0.093	0.000	0.000	0.000	1.000	112,633
Car Service	0.160	0.297	0.000	0.000	0.214	2.000	112,633
Interstate	0.004	0.049	0.000	0.000	0.000	1.000	112,633

Finally, we provide summary statistics for the observation-level data in our analysis in Table 9. The greatest number of stations a brand has in a single county in our data is 83. The 25 percentile is 2, and we have several observations of a brand with only a single station in our market. The variable Wholesale FE is the average wholesale price for a brand within a county. We interact this variable with the U.S. oil production data to generate an instrument for price in the demand estimation.

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.

4.3 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 free-agent type customers. We then estimate demand using the typical logit demand regression. Our chosen dynamic parameters minimize the

Our regression equation takes the following form:

$$\ln(s_{jmt}(0)/s_{0mt}(0)) = \xi_{jM} - \alpha p_{jmt} + \gamma N_{jmt} + X_{jmt}\beta + \eta_{jmt}$$

Here, we have added the subscript m to denote the market (county) and M to denote the larger region (state). We have shares and prices at the brand-county-week level. We model a within county share as depending on the number of stations for the brand in that market, N_{jmt}, after accounting for a brand-state fixed effect, ξ_{jM}. We allow for endogeneity in pricing behavior by instrumenting for p_{imt} with predicted deviations in wholesale costs, where the predictions are obtained from a regression of deviations of wholesale costs (from the brand-state average) on the interaction of US production of crude oil (obtained from EIA) with the average wholesale cost for the brand in the state.[19] This gives us brand-state-specific time variation in our instrument, and it is plausibly tied to variation in the wholesale cost and not linked to demand. We chose this measure, rather than instrumenting directly with brand-state wholesale costs, to account for the possibility that wholesalers respond to local, brand-specific demand shocks.

[19] Our measure of the average brand-state wholesale cost is the fixed effect obtained by a regression of wholesale costs on brand state and weekly fixed effects, thereby accounting for compositional differences in brand-states across time.

Figure 4: Shares and Prices

(a) Total Market Shares and Prices

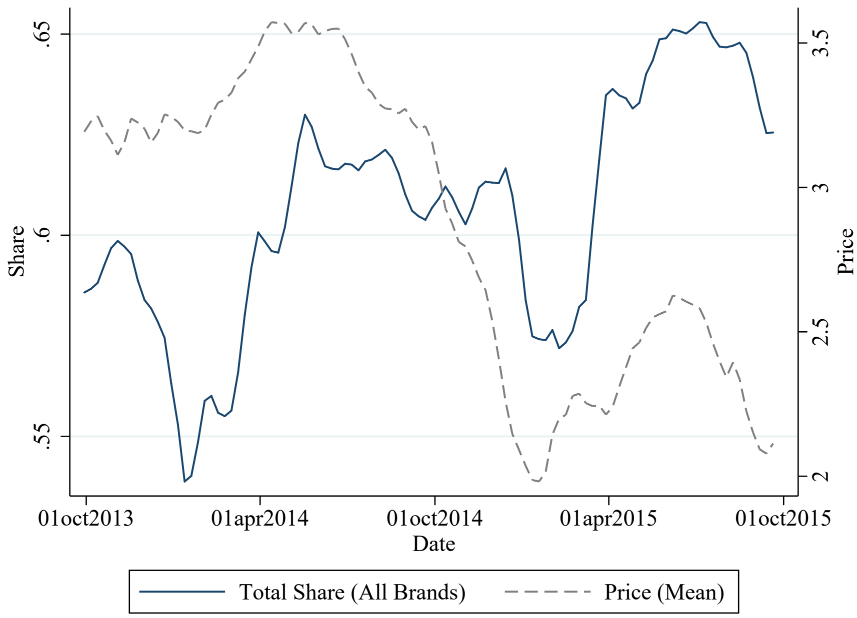

(b) Instrument and Prices

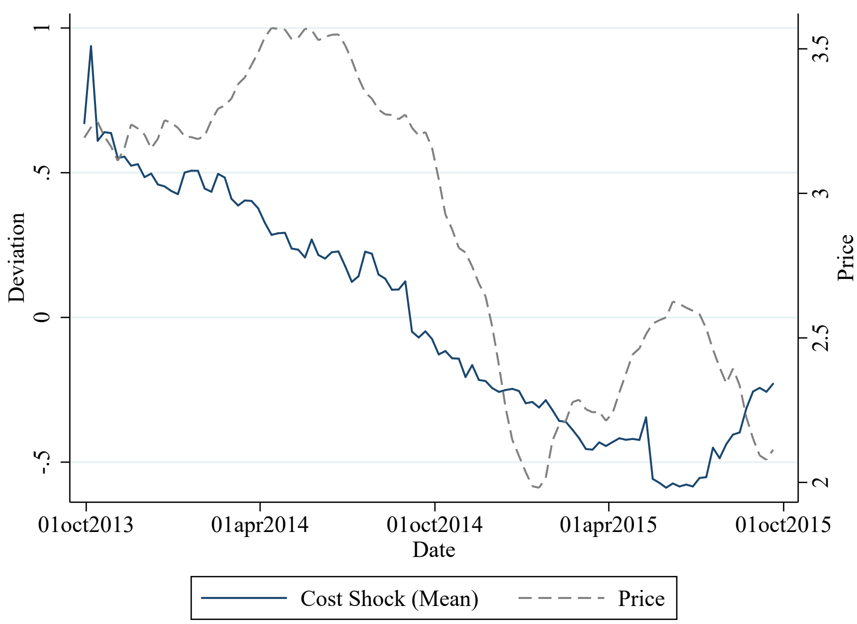

Table 10: Demand Regressions

	Static Model			Dynamic Model
	(1)	(2)	(3)	(4)
Price	−0.009	−0.234***	−2.536***	−3.019***
	(0.013)	(0.015)	(0.318)	(0.505)
Number of Stations	0.019***	0.019***	0.072***	0.081***
	(0.004)	(0.004)	(0.011)	(0.013)
Food	0.134***	0.128***	0.137***	0.163***
	(0.022)	(0.022)	(0.031)	(0.040)
Supermarket	−0.013	−0.013	0.020	0.026
	(0.021)	(0.020)	(0.025)	(0.031)
Car Service	0.002	−0.000	0.006	0.003
	(0.024)	(0.024)	(0.025)	(0.032)
Interstate	0.026	0.026	0.037***	0.046***
	(0.017)	(0.017)	(0.014)	(0.016)
Dynamic Parameters				
δ				0.762
				(0.007)
$\bar{\xi}$				2.180
				(0.061)
$\bar{\alpha}$				2.537
				(0.056)
IV	No	Yes	Yes	Yes
Demographic Controls	X	X	X	X
Week FEs			X	X
County-(Week of Year) FEs			X	X
Brand-State FEs			X	X
Observations	112,417	112,417	112,417	112,417
R^2	0.066	0.045	0.441	0.199

Notes: Significance levels: * 10 percent, ** 5 percent, *** 1 percent. The 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 free agent customers, which are calculated based on the estimated dynamic parameters. Standard errors are clustered at the county level. Standard errors for the dynamic parameters (are preliminary and) are calculated via the bootstrap.

In addition to the instrument, we employ panel data methods to address other unobservables. We allow for the fact that η_{jmt} may be correlated over time. We let $\eta_{jmt} = \phi_t + \psi_{mw(t)} + \kappa_{jmt}$, and estimate weekly fixed effects, $\{\phi_t\}$, along with county-specific (weekly) seasonal demand shocks, $\{\psi_{mw(t)}\}$. We also estimate brand-state specific fixed effects, $\{\xi_{jM}\}$. Once we incorporate these fixed effects, the exclusion restriction for a valid use of our instrument is that the brand-market-period specific shock κ_{jmt} is uncorrelated with the instrument, after accounting for aggregate period-specific shocks and brand-state level differences.[20]

The estimates are reported in Table 10. The first three columns report coefficient estimates using the observed shares in logit demand estimation. The fourth column reports the results from our dynamic model, along with the dynamic parameters. In addition to brand amenities, we also control for the following demographic characteristics: median household income, population, commute percent, and population density. Each of these variables are weighted by the number of stations in each census tract within a county. Thus, entry and exit by a brand give us variation in these characteristics over time.

We estimate that 76 percent of customers become locked-in week-over-week after purchasing from a particular station. These customers have low price sensitivity, as the net price sensitivity $\bar{\alpha} + \alpha$ of approximately -0.5 is quite small. The price coefficient, -3.02, is greater than that of the static model (-2.54).

Table 11: Implied Elasticities from Dynamic Model

Group	Mean	Pctl(25)	Median	Pctl(75)
Free Agent	-8.099	-9.447	-8.246	-6.692
Locked-In	-0.028	-0.009	-0.002	-0.001
Weighted	-1.975	-2.501	-1.842	-1.263
Naive (Static)	-6.251	-7.444	-6.238	-5.156

To interpet the price coefficients, we summarize the implied elasticities in Table 11. The locked-in customers of our model are very inelastic, with a near-zero response to price effects. The free agents, however, are highly elastic, with an average own-price elasticity of -8.1. This is large in magnitude, and it implies that for a 1 percent increase in price (roughly 3 cents), the station will lose 8.1 percent of the free-agent customers. This high level of price sensitivity for a subset of retail gasoline customers seems plausible, as, anecdotally, some “shoppers” are known to go well out of the way to save a few cents per gallon.

The average weighted elasticity, which weighs free-agent and locked-in customers by their relative (purchasing) proportions, is -2.0. This weighted elasticity is starkly different than the elasticity one obtains by estimating the static model. A "naive" estimate (supposing

[20] An alternative interpretation of our decomposition is that we attribute all of the brand-specific correlation in demand over time within a market to unobservable demand types arising from consumer loyalty.

the true model were dynamic) of this elasticity would obtain a value of -6.3, which implies a much greater loss in market share for a given price increase than we obtain from the dynamic model. Indeed, the static model obtains an elasticity more in line with the free-agent elasticity than the overall elasticity from the dynamic model.

Additionally, we find that variables we would expect to increase demand, including the number of stations a brand has in a market and the availability of amenities such as food, a supermarket, or proximity to an interstate, are positively correlated with utility.

Table 12: Shares of Locked-In and Free-Agent Customers

Statistic	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Observed Share (S_{jmt})	0.141	0.110	0.000	0.061	0.187	0.688
Free-Agent Share ($s_{jmt}(0)$)	0.065	0.063	0.000	0.022	0.089	0.686
Locked-In Share ($s_{jmt}(j)$)	0.979	0.073	0.039	0.992	1.000	1.000
r_{jmt}	0.107	0.084	0.000	0.046	0.142	0.524
Portion from Free Agents	0.244	0.141	0.000	0.166	0.309	1.000
Portion from Locked-In	0.753	0.145	0.000	0.689	0.833	1.000

For ease of interpretation of the dynamic parameters, we present summary statistics and graphs of the resulting shares. Table 12 displays the means and standard deviations for observed shares and the two components of the shares we identify. We also report the portion of the observed share that a firm realizes from the locked-in customers, which is 75 percent on average. Firms (brands) retain nearly all of their locked-in customers (98 percent), and capture 6.5 percent of the free-agent customers on average.

In Figure 5, we display the time series of the average price and the average shares. The shares for the customers that are not locked-in, displayed in Panel (a), are negatively correlated with the price.[21] Intuitively, customers shift to the outside option when prices are higher. On the other hand, the share of locked-in customers, displayed in Panel (b), track the prices more closely. This occurs because prices respond positively to demand shocks, as do the locked-in customers, who we find to be price inelastic. Figure 6 displays the elasticities over time for each type of customer. The weighted elasticity, displayed in Panel (b), as well as the free-agent elasticity, are negatively correlated with prices. The locked-in elasticity is relatively constant.

Figure 7 presents the average portion of customers for each firm that are coming from the locked-in set of customers. In Panel (b), we compare the portions for large brands and fringe brands (defined as both the synthetic FRINGE brand and unbranded stations), as we see that the branded stations get a relatively higher portion of their shares from locked-in customers. This is consistent with a story of brand loyalty.

[21] For the remainder of the figures in this section, including this one, our plots display equally-weighted four-week moving averages.

Finally, we present two further cuts of the data between branded firms and fringe firms. Figure 8 presents the shares for each type of customer over time. Branded and fringe firms attract approximately the same amount of free-agent customers, but branded firms have higher shares of the locked-in customers, consistent with the above finding. Correspondingly, branded firms have consistently less elastic locked-in customers, as displayed in Figure 9.

Figure 5: Shares Over Time

(a) Free-Agent Customers

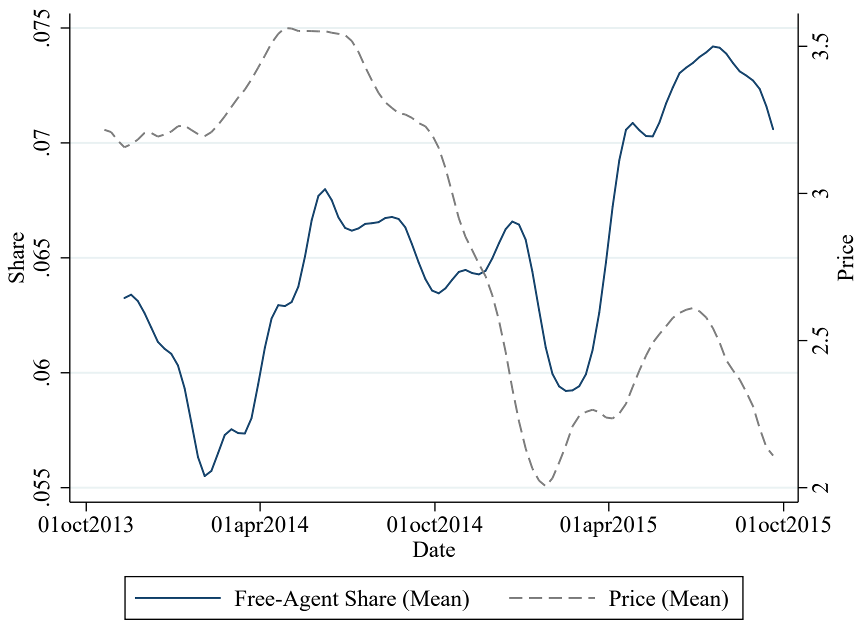

(b) Locked-In Customers

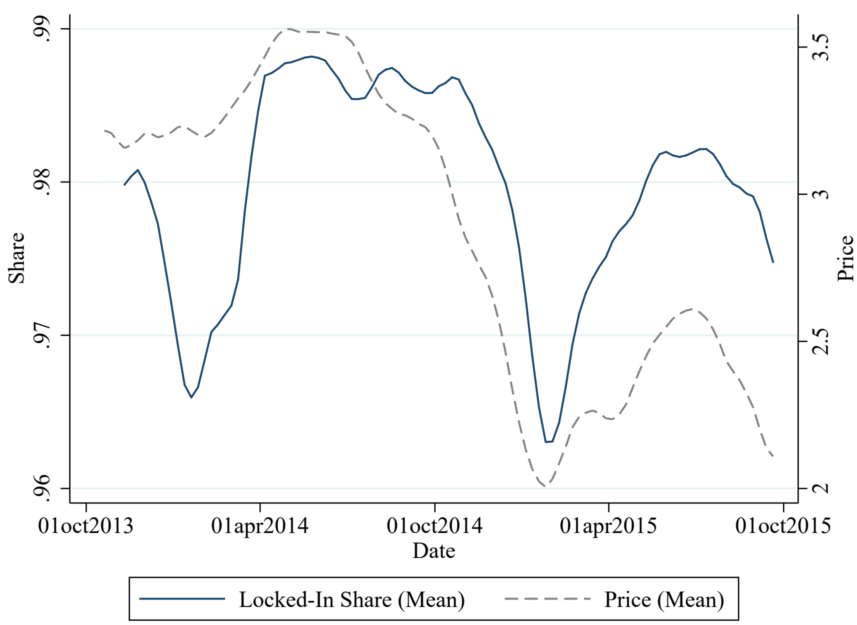

Figure 6: Elasticities Over Time

(a) Free Agent and Locked-In

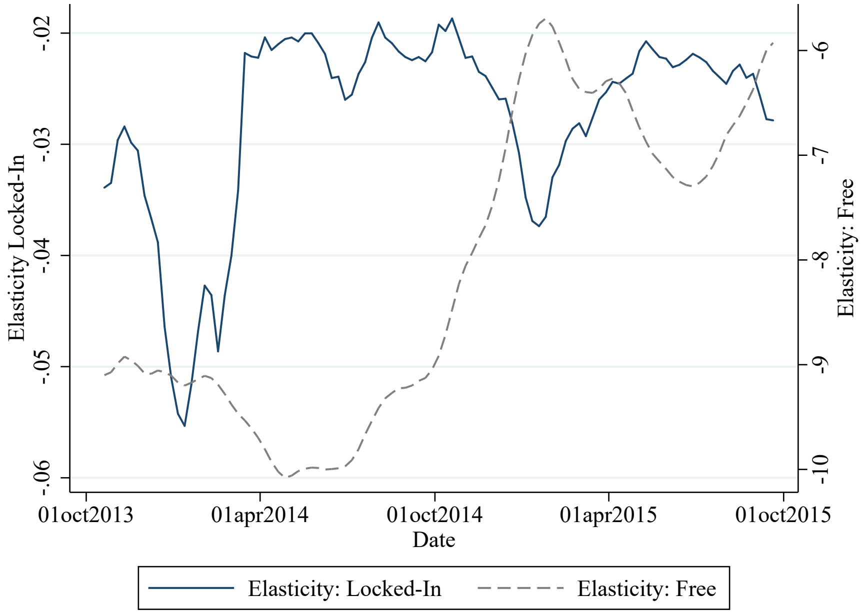

(b) Mean Elasticity

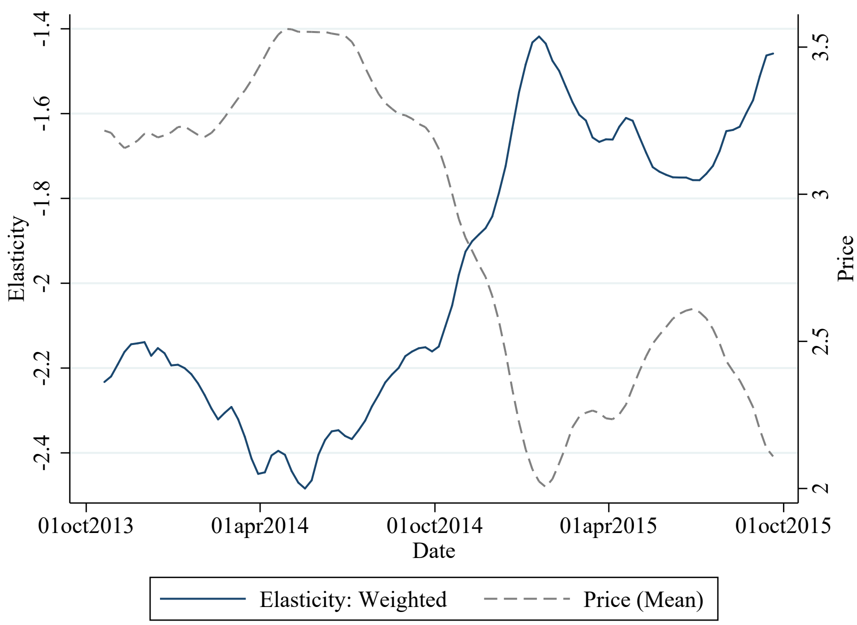

Figure 7: Portion from Locked-In Customers

(a) All Customers and Shares

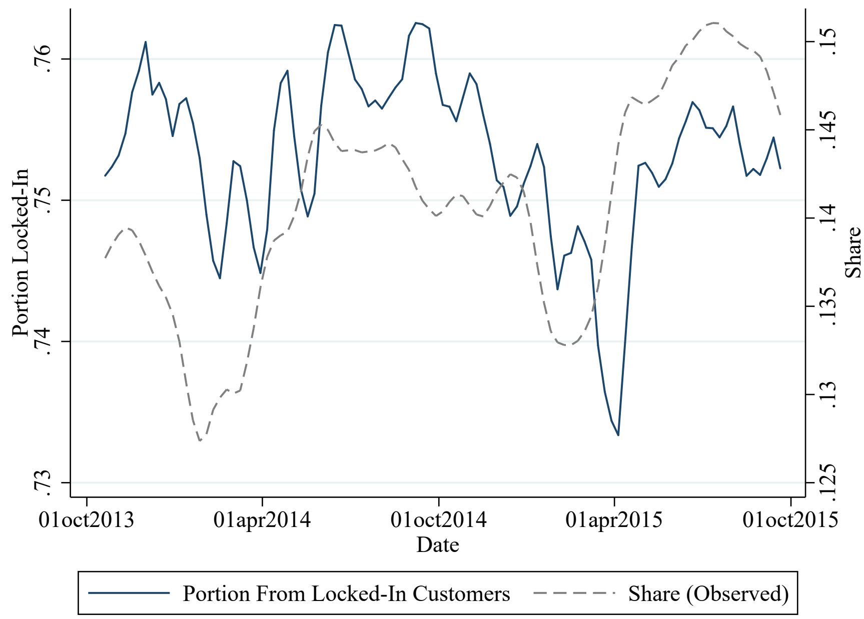

(b) Large Brands versus Fringe

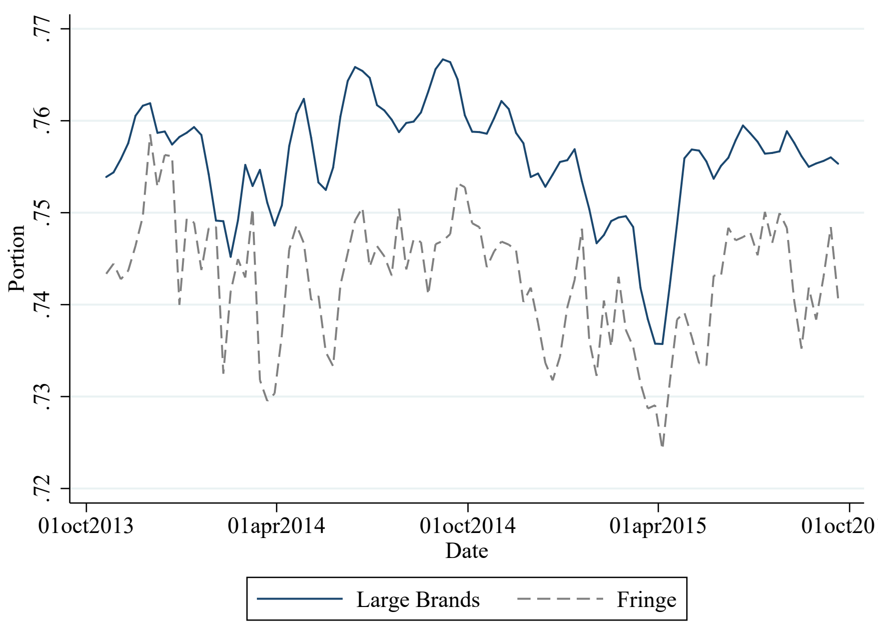

Figure 8: Shares Over Time

(a) Free-Agent Customers

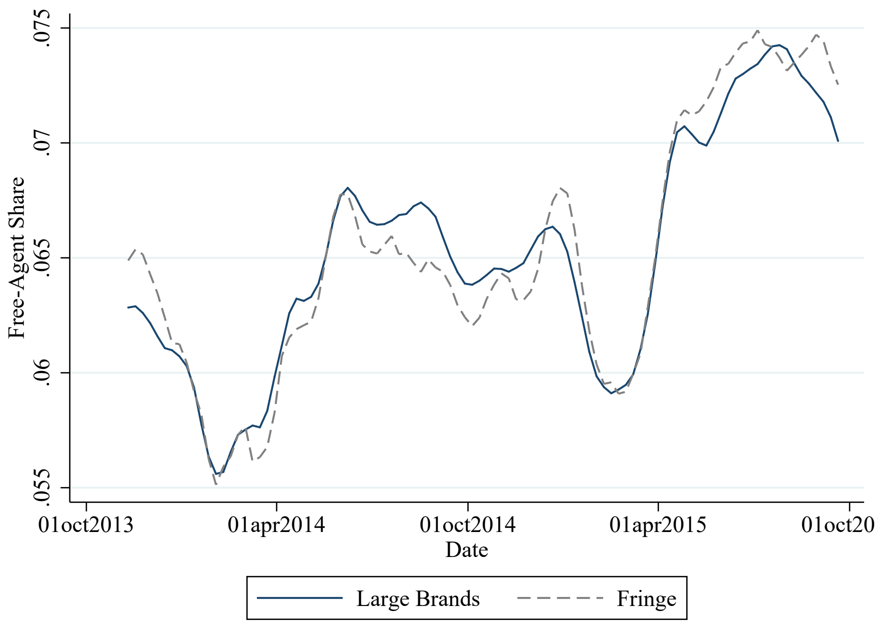

(b) Locked-In Customers

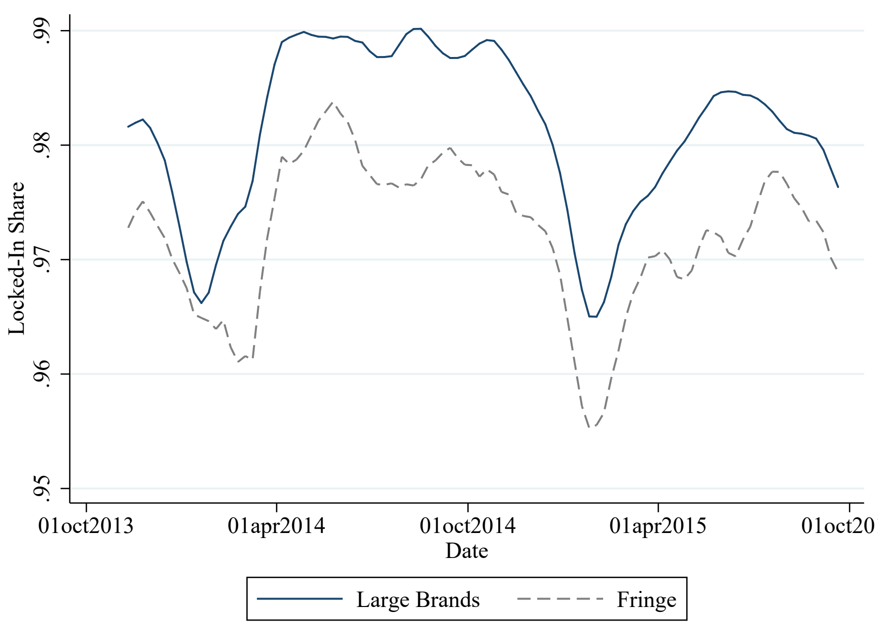

Figure 9: Elasticities: Large Brands versus Fringe

(a) Locked-In Elasticity

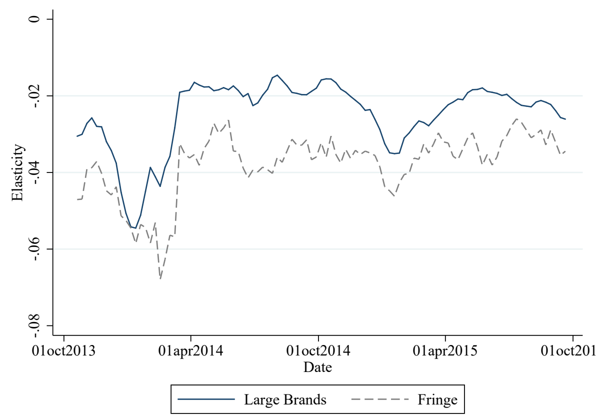

(b) Free-Agent Elasticity

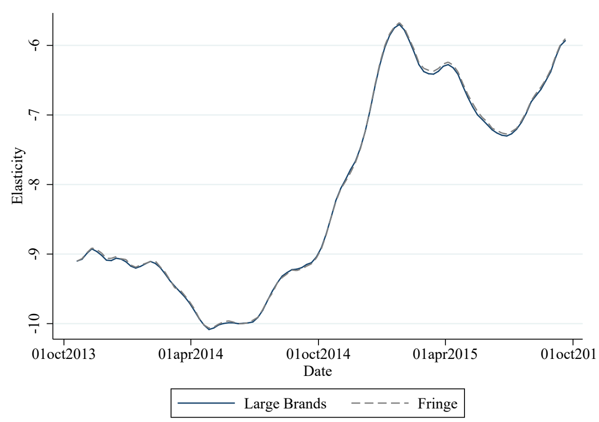

(c) Weighted Elasticity

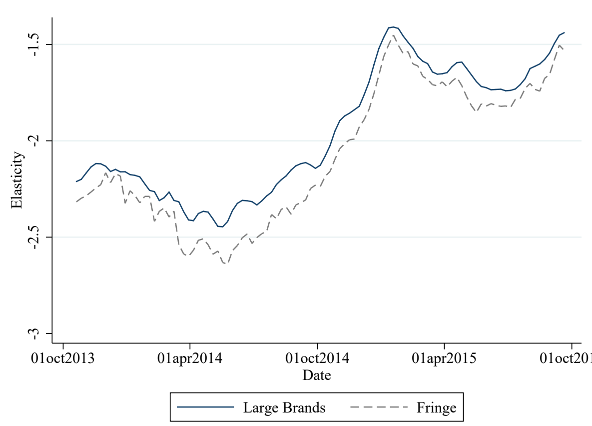

5 Empirical Application: Supply-Side Analysis

5.1 Supply-Side Estimation

Given our demand estimates, we construct the components in each firm's Bellman equation from (5). This allows us to directly estimate the derivative of the continuation value, $\beta E[V(\boldsymbol{r}', \boldsymbol{c}')|\boldsymbol{r}, \boldsymbol{c}, \boldsymbol{p}]$. We use the expected and unexpected costs estimated in our pass-through section to estimate a model-restricted function for the continuation value. This estimated function allows us to pursue a counterfactual merger simulation, which we do in the next section.

Using our estimated demand parameters, we are able to recover the (derivative of) the continuation value. Recall the dynamic FOC for the firm:

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

We include state variables x that factor into firm beliefs and expectations but do not affect demand. Our demand-side estimation allows for a direct estimate of the static FOC: $\frac{\partial S_{jt}}{\partial p_{jt}}\left(p_{jt}-c_{jt}\right)+S_{jt}$. The implied derivative of the continuation value is the negative value of the static FOC, which rationalizes the observed price given Bertrand price competition.

5.2 Results: Supply-Side Behavior

Summary statistics for the value of the derivative of the continuation value (DCV) are presented in Table 13. The mean and median are negative, which means that firms price lower than the static profit-maximizing price implied by our demand model. The magnitudes are significant: as the average (scaled) profit in our data is 0.029, this implies a roughly 4 percent increase in profits from raising prices by 1 cent (recall that margins are approximately 21 cents per gallon).

We interpret this as arising from a dynamic incentive, though other explanations may be plausible.[22] Intuitively, firms are lowering prices to invest in future demand.

Table 13: Summary of Implied $\beta\frac{\partial E[V_j(\cdot)|\cdot]}{\partial p_{jt}}$

	mean	p25	p50	p75
All	-0.122	-0.162	-0.093	-0.051
KY	-0.141	-0.186	-0.111	-0.062
VA	-0.107	-0.143	-0.082	-0.045

[22]For example, a component of this residual may be profits obtained by complementary products, such as food sold at retail gasoline stations.

To generate counterfactuals, we estimate how the DCV changes with state variables and prices. We project the DCV onto several variables, including market characteristics, demand shocks, expected and unexpected costs, and prices. The results are reported in Table 14. The first two specifications report the coefficients from a regression. Specifications (3) and (4) report the net effect once we run a regression with interactions for all of the variables.

Table 14: Dynamic FOC Regressions: $dE[V]/dP$

	(1) Main	(2) Main	(3) Interactions	(4) Interactions
Cost, Exp.	-0.098^{***}	-0.111^{***}	-0.098^{***}	-0.098^{***}
	(0.001)	(0.015)	(0.000)	(0.005)
Cost, Unexp.	-0.098^{***}	-0.096^{***}	-0.098^{***}	-0.101^{***}
	(0.001)	(0.003)	(0.000)	(0.001)
Cost, Tax	-0.093^{***}	-0.057^{***}	-0.093^{***}	-0.092^{***}
	(0.001)	(0.003)	(0.000)	(0.001)
Locked-in Portion	-1.136^{***}	-0.873^{***}	-1.136^{***}	-1.062^{***}
	(0.002)	(0.004)	(0.001)	(0.001)
Price	0.099^{***}	0.110^{***}	0.099^{***}	0.101^{***}
	(0.002)	(0.003)	(0.001)	(0.001)
Price Squared	0.001^{***}	-0.001^{**}	0.001^{***}	0.001^{***}
	(0.000)	(0.000)	(0.000)	(0.000)
Cost, Exp. Change (30 d)	0.000	-0.012	0.000	0.000
	(0.001)	(0.015)	(0.000)	(0.005)
Demand Shock	-0.003^{***}	-0.003^{***}	-0.003^{***}	-0.005^{***}
	(0.000)	(0.000)	(0.000)	(0.000)
Mean Price (Rivals)	0.014^{***}	0.011^{***}	0.014^{***}	0.013^{***}
	(0.001)	(0.002)	(0.000)	(0.001)
S.D. Price (Rivals)	0.004^{**}	0.002	0.004^{***}	0.003^{***}
	(0.002)	(0.002)	(0.001)	(0.001)
County-Brand + Time FEs		X		X
Structure Controls	X	X	X	X
Observations	109191	109191	109191	109191
R^2	0.96	0.97	0.99	1.00

Standard errors in parentheses

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

As the derivative of the continuation value with respect to price is negative, negative coefficients imply a greater negative marginal impact of price on the continuation value. Thus, negative coefficients, which move the DCV further from zero, may be interpreted as exacerbating the effect of dynamics on pricing behavior. For example, the greater the amount of locked-in customers a brand has, the greater the dynamic effect on the pricing decision.

Table 15: Dynamic FOC Regressions: Sensitivity [$\ln(|dE[V]/dP|)$]

	(1) Main	(2) Main	(3) Interactions	(4) Interactions
Cost, Exp.	0.881^{***}	1.572^{***}	0.958^{***}	1.428^{***}
	(0.014)	(0.211)	(0.014)	(0.200)
Cost, Unexp.	0.797^{***}	0.713^{***}	0.870^{***}	0.845^{***}
	(0.016)	(0.038)	(0.016)	(0.036)
Cost, Tax	1.056^{***}	0.615^{***}	1.123^{***}	0.847^{***}
	(0.020)	(0.042)	(0.019)	(0.039)
Locked-in Portion	5.733^{***}	5.707^{***}	6.033^{***}	7.418^{***}
	(0.043)	(0.061)	(0.042)	(0.058)
Price	-1.238^{***}	-1.069^{***}	-1.290^{***}	-0.939^{***}
	(0.032)	(0.041)	(0.031)	(0.039)
Price Squared	-0.018^{***}	0.012^{*}	-0.016^{***}	-0.016^{**}
	(0.005)	(0.007)	(0.005)	(0.006)
Cost, Exp. Change (30 d)	0.094^{***}	0.961^{***}	0.101^{***}	0.586^{***}
	(0.013)	(0.223)	(0.012)	(0.210)
Demand Shock	0.170^{***}	0.053^{***}	0.160^{***}	0.034^{***}
	(0.002)	(0.002)	(0.002)	(0.002)
Mean Price (Rivals)	0.187^{***}	0.105^{***}	0.243^{***}	0.173^{***}
	(0.026)	(0.024)	(0.025)	(0.022)
S.D. Price (Rivals)	0.277^{***}	-0.127^{***}	0.264^{***}	-0.168^{***}
	(0.034)	(0.030)	(0.033)	(0.028)
County-Brand + Time FEs		X		X
Structure Controls	X	X	X	X
Observations	109191	109191	109191	109191
R^2	0.77	0.90	0.78	0.91

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

To show more directly how sensitive firms are to dynamic considerations, we also report the results from regressions where we replace the value of the DCV with the logged absolute value. This provides a measure of sensitivity, and larger coefficients indicate that dynamic considerations are more impactful. The results are reported in Table 15. After accounting for interactions among the variables, we find that changes in unexpected costs have a greater impact on the continuation value than changes in expected costs. This corroborates the reduced-form pass-through results from Figure 3, in which we find that the pass-through for unexpected costs is lower than for expected costs. Our model rationalizes this phenomenon as arising from a change in the perceived impact of price on the continuation value.

5.3 Merger Simulation

To evaluate the effect of the dynamic effects in a merger context, we simulate a merger between Marathon and BP, which are the number one and number four (non-fringe) in terms of overall shares in our sample. Out of the 252 markets, they overlap in 74. In these 74 markets, the average HHI is 1500,[23] and the mean change in HHI resulting from the merger is 387. In 8 markets, the resulting HHIs are greater than 2500 and the changes are greater than 200, which meet the typical thresholds that are presumed likely to enhance market power. We allow the firms to merge at the beginning of September 2014, and we calculate counterfactual prices and shares for the second half of our sample.

Our setting, with dynamic brand effects, puts a greater emphasis on the what exactly a merger will be in practice. This is a nuanced question. Do the merging firms retain separate brands, or convert all stations to a single brand? Do the customers of the old brand become immediately attached to the new, or is there no transfer in brand loyalty?

Table 16: Merger Effects: Brand Loyalty Tranfers

	Brands	Price	Share	Profit
1	Marathon-BP	0.81	-0.20	7.70
2	Other	-0.11	0.33	-2.80
3	Overall	0.29	-0.01	2.52

We consider two scenarios. In the first scenario, the merging firms retain all brand loyal customers from both Marathon and BP. Table 16 displays the mean effects when all brand effects transfer. The effects are quite small, with an average price increase for the merging firms of less than one percent. Shares for these firms barely decline. As this is a low-margin industry, these effects have an economically meaningful effect on profits, which increase by 7.7 percent. Profits from the competitors fall, which results in almost no net effect on industry profits. Interestingly, prices for competitors fall, which is the opposite of what we usually expect after a merger.

In our second scenario, brand loyalty does not transfer at all: all of the BP customers become free agents. Table 2 displays the results of this scenario. As expected, the loss of brand-loyal customers decreases profits for the merging firms relative to the first scenario. However, the merger occurs over a long enough time frame that most of this effect appears to be compensated for by long-run behavior. Overall prices rise by about the same amount,

[23]As we treat the fringe as a profit-maximizing entity, we calculate HHI treating the fringe as one firm. This will overstate the baseline HHI and the price response by competitors. The change in HHI is unaffected by this abstraction.

Table 17: Merger Effects: Brand Loyalty Does Not Transfer

	Brands	Price	Share	Profit
1	Marathon-BP	0.81	-0.76	6.79
2	Other	-0.11	0.44	-2.69
3	Overall	0.28	-0.20	2.09

but overall shares fall more, relative to the first scenario. Intuitively, the stock of brand loyal customers keep quantities up even when prices are raised.

Table 18: Merger Effects: Regressions

	Price		Share		Profits	
	(1)	(2)	(3)	(4)	(5)	(6)
HHI Delta (100)	0.129^{***}		-0.059^{*}		1.072^{***}	
	(0.024)		(0.033)		(0.271)	
Combined r		3.273^{***}		-1.295^{*}		29.451^{***}
		(0.493)		(0.725)		(5.503)
Observations	74	74	74	74	74	74
R^2	0.274	0.376	0.040	0.042	0.176	0.282

Notes: Significance levels: * 10 percent, ** 5 percent, *** 1 percent. The table displays regressions of the estimated merger effects for the merging brands on the change in HHI resulting from the merger and the combined locked-in customers entering the merger. No constant is used in the regressions.

In Table 18, we regress the percentage effects for each market on indicators of the strength of the merger: the HHI delta and the combined level of locked-in customers for the merging brands. The coefficient on the first model indicates that an HHI change of 100 points would result, on average, in a 0.1 percent price change. For the second model, the coefficient of 3.3 indicates that a combined loyalty share of 0.3 would result in a 1 percent change in prices.

Table 19: Merger Effects: Static Model

	Brands	Price	Share	Profit
1	Marathon-BP	4.69	-17.22	31.45
2	Other	0.65	5.70	18.04
3	Overall	2.33	-5.68	24.33

Overall, these effects are small relative to what one might normally expect from a merger. Even for the 8 markets that fall under the category of likely to increase market power, the results are quite similar. For comparison, we report the results from a merger analysis using a static model in Table 19, which has average price effects of approximately 5 percent. How do we interpret dynamic model? Though the merged firm has a static incentive to raise prices, the dynamic incentive is also strengthened, as the firm accumulates more valuable dynamic assets. These effects work to offset each other, relative to the effects arising in a static analysis.

6 Conclusion

We develop a model of dynamic demand that accounts for the slow adjustment of prices to changes in cost. The dynamics result from competing firms optimally setting prices to customers that may become loyal to their current supplier. Using data from retail gasoline markets, we first demonstrate that prices adjust slowly to cost changes, and that path of price adjustment depends upon whether the costs change is expected or unexpected. This finding demonstrates that importance of accounting for firm expectations when estimating pass-through and estimating a demand model that can accommodate these short-run dynamics.

We derive a new estimator that can identify dynamic demand parameters using data on price, shares, and an instrument. Preliminary results suggest that about three-fourths of retail gasoline customers become locked in to the firm from which they currently purchase on a week-to-week basis, and that these loyal consumers are extremely price insensitive. Conversely, we find that non-loyal customers are quite price sensitive. We evaluate the dynamic incentives affecting prices, and we show, both theoretically and empirically, that merger effects are muted by the presence of dynamic incentives.

References

Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating Dynamic Models of Imperfect Competition. *Econometrica 75*(5), 1331–1370.

Beggs, A. and P. Klemperer (1992, May). Multi-Period Competition with Switching Costs. *Econometrica 60*(3), 651.

Bergemann, D. and J. Välimäki (2006, March). Dynamic Price Competition. *Journal of Economic Theory 127*(1), 232–263.

Borenstein, S., A. C. Cameron, and R. J. Gilbert (1997). Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes? *Quarterly Journal of Economics 112*(1), 305–39.

Chandra, A. and M. Tappata (2011). Consumer Search and Dynamic Price Dispersion: An Application to Gasoline Markets. *The RAND Journal of Economics 42*(4), 681–704.

Dube, J. P., G. Hitsch, and P. Rossi (2009). Do Switching Costs Make Markets Less Competitive? *Journal of Marketing Research 46*(4), 11.

Eckert, A. (2013). Empirical Studies of Gasoline Retailing: A Guide to the Literature. *Journal of Economic Surveys 27*(1), 140–166.

Engle, R. F. and C. W. J. Granger (1987, March). Co-Integration and Error Correction: Representation, Estimation, and Testing. *Econometrica 55*(2), 251.

Farrell, J. and P. Klemperer (2007). Chapter 31 Coordination and Lock-In: Competition with Switching Costs and Network Effects. In *Handbook of Industrial Organization*, Volume 3, pp. 1967–2072. Elsevier.

Farrell, J. and C. Shapiro (1988). Dynamic Competition with Switching Costs. *The RAND Journal of Economics 19*(1), 123.

Hendel, I. and A. Nevo (2013, December). Intertemporal Price Discrimination in Storable Goods Markets. *American Economic Review 103*(7), 2722–2751.

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*(2), 672–682.

Miller, N. H., M. Remer, and G. Sheu (2013, March). Using Cost Pass-through to Calibrate Demand. *Economics Letters 118*(3), 451–454.

Neumark, D. and S. A. Sharpe (1992, May). Market Structure and the Nature of Price Rigidity: Evidence from the Market for Consumer Deposits. *The Quarterly Journal of Economics 107*(2), 657–680.

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*(2), 373–399.

Peltzman, S. (2000, June). Prices Rise Faster than They Fall. *Journal of Political Economy 108*(3), 466–502.

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

Slade, M. E. (1998). Optimal Pricing with Costly Adjustment: Evidence from Retail-Grocery Prices. *The Review of Economic Studies 65*(1), 87–107.

A Duopoly Steady-State Analysis

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

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

2. Solve for p^{ss} and r^{ss} using the steady-state restrictions.

3. Take the numerical derivative of p^{ss} with respect to r^{ss}. We approximate the numerical derivative by slightly perturbing r^{ss} by h, resolving for p^{ss}, and calculating $\frac{p^{ss}(r+h)-p^{ss}(r-h)}{2h}$.

4. Solve again for p^{ss} and r^{ss} using the steady-state restrictions and the updated $\frac{dp}{dr}$.

5. If the changes in p^{ss} and r^{ss} between steps 2 and 4 fall below a critical value, and the changes in each of the four elements of $\frac{dp}{dr}$ between steps 1 and 3 fall below a critical then we have found all steady-state values. If the change in any of the unknowns is above the critical value then repeat steps 1-4 using the updated values of $\frac{dp}{dr}$.

B Common and Idiosyncratic Costs

In Figure 10, we present Pass-through results for common market-level costs. In Figure 11, we present Pass-through results controlling for the other firm's prices.

Figure 10: Cumulative Pass-through for Common Costs

(a) Unexpected Costs

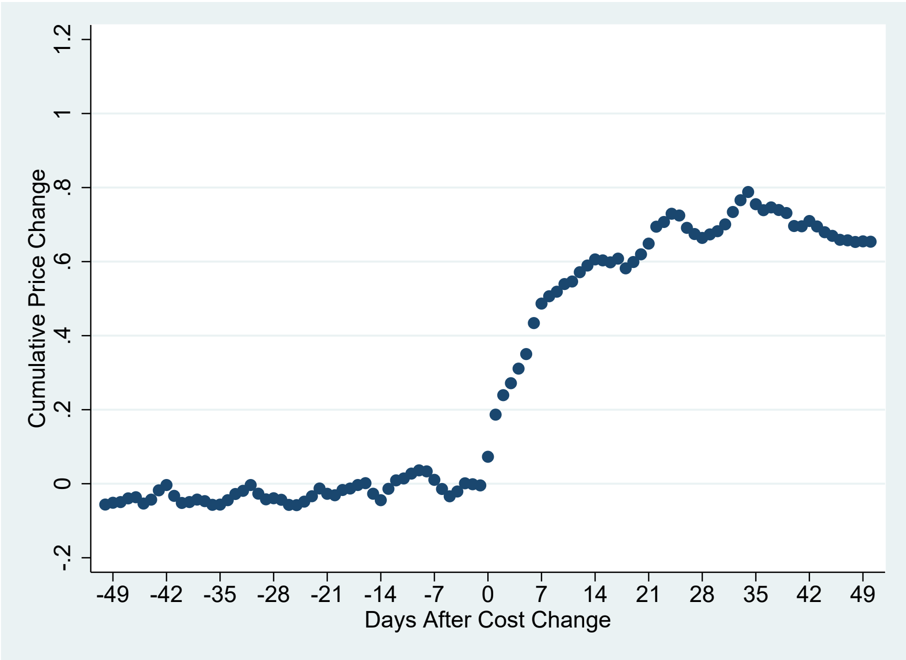

(b) Expected Costs

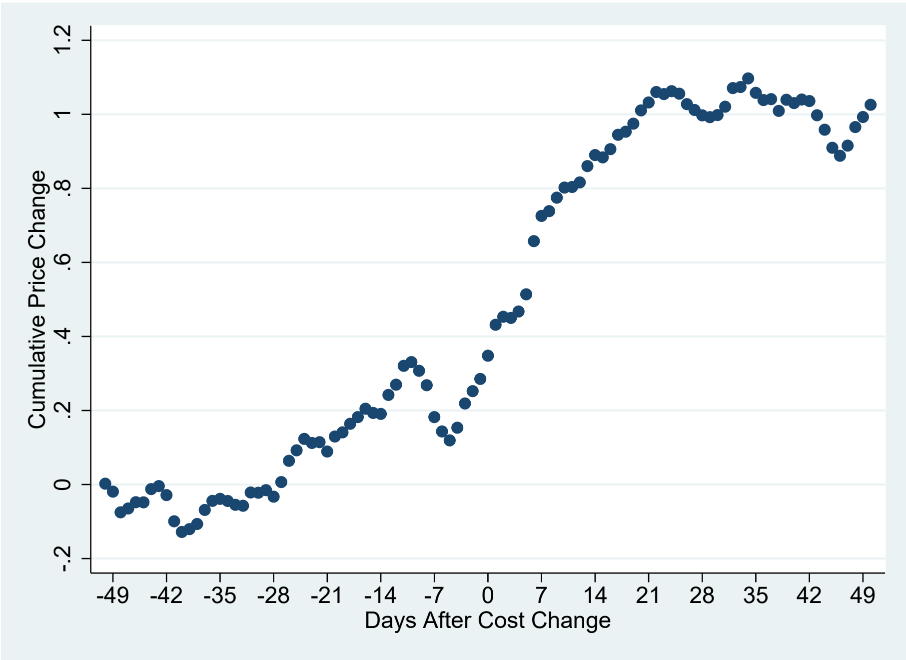

Figure 11: Cumulative Pass-through Controlling for Rival Prices

(a) Unexpected Costs

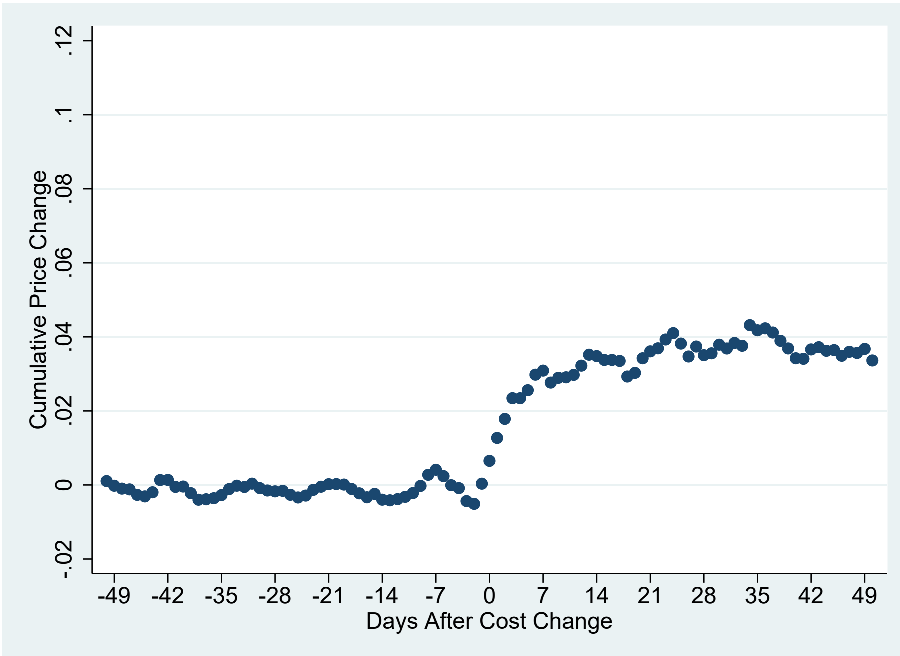

(b) Expected Costs

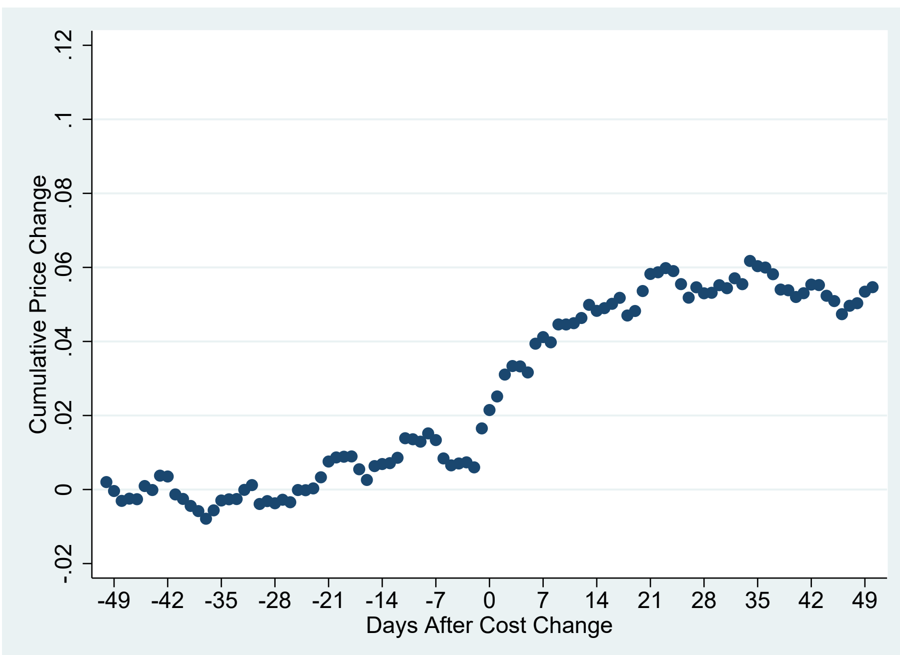

C Four-Parameter Heterogeneity Specification

Here we report summary statistics four our four-parameter estimates of firm-specific heterogeneity.

Table 20: Pass-through Heterogeneity: Four-Parameter Specification

Parameter	p10	p25	p50	p75	p90
Anticipation	-23.8	-17.2	-6.5	-2.2	0.4
Rate	0.010	0.016	0.025	0.039	0.066
Duration	9.8	16.5	29.0	48.2	60.8
Finish	6.9	12.7	21.9	32.3	40.0
LR_PTR	0.362	0.532	0.711	0.896	1.177
Anticipation_Exp	-43.7	-36.3	-28.2	-16.5	-5.4
Rate_Exp	0.015	0.018	0.022	0.029	0.047
Duration_Exp	21.6	35.3	51.7	63.4	73.5
Finish_Exp	10.1	15.6	22.8	29.6	37.0
LR_PTR_Exp	0.883	0.980	1.089	1.219	1.340
Offset	-0.16	-0.11	-0.06	-0.00	0.05
Offset_Exp	-0.27	-0.17	-0.07	0.02	0.10
Above_Zero	0.261	0.448	0.655	0.863	1.155
Above_Zero_Exp	0.855	0.937	1.014	1.102	1.215